arXiv:1407.5234v1 [cs.IT] 20 Jul 2014

Compressed Subspace Matching on the Continuum

William Mantzel and Justin Romberg*

July 22, 2014

Abstract

We consider the general problem of matching a subspace to a signal in $\mathbb{R}^N$ that has been observed indirectly (compressed) through a random projection. We are interested in the case where the collection of K-dimensional subspaces is continuously parameterized, i.e. naturally indexed by an interval from the real line, or more generally a region of $\mathbb{R}^D$. Our main results show that if the dimension of the random projection is on the order of K times a geometrical constant that describes the complexity of the collection, then the match obtained from the compressed observation is nearly as good as one obtained from a full observation of the signal. We give multiple concrete examples of collections of subspaces for which this geometrical constant can be estimated, and discuss the relevance of the results to the general problems of template matching and source localization.

1 Introduction

We consider the general problem of finding the subspace that best approximates a fixed signal $\boldsymbol{h}_0 \in \mathbb{R}^N$ from a parameterized collection of K-dimensional subspaces $\{\mathcal{S}_\theta : \theta \in \Theta\}$, where θ is a parameter vector chosen from some compact parameter set $\Theta \subset \mathbb{R}^D$. Given this collection, the subspace best matched to $\boldsymbol{h}_0$ is the solution to

$$\bar{\theta} = \arg\min_{\theta\in\Theta} \min_{\boldsymbol{h}\in\mathcal{S}_\theta} \|\boldsymbol{h}_0 - \boldsymbol{h}\|_2^2 = \arg\min_{\theta\in\Theta} \|\boldsymbol{h}_0 - \boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2. \tag{1}$$

The operator $\boldsymbol{P}_\theta$ above is the orthogonal projection onto subspace $\mathcal{S}_\theta$. In words, program (1) returns the (index of the) subspace in the collection which contains the closest point to $\boldsymbol{h}_0$ in terms of the standard Euclidean distance.

In this paper, we explore how effectively this matching can be done from a set of indirect linear observations of $\boldsymbol{h}_0$. In particular, we will quantify how effectively this problem can be solved when our observations $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{h}_0$ are *compressed samples*, meaning $\boldsymbol{\Phi}$ is a $M \times N$ *underdetermined* matrix whose rows are diverse. For simplicity, we will consider the case where the entries of $\boldsymbol{\Phi}$ are independent and Gaussian, but we expect that the majority of our results could be extended to other types of measurement scenarios.

*W.M. is at Qualcomm in San Diego, CA; J.R. is in School of Electrical and Computer Engineering at Georgia Tech in Atlanta, GA. Email: wmantzel@gmail.com, jrom@ece.gatech.edu. This work was supported by ONR grant N00014-11-1-0459 and a grant from the Packard Foundation.

Given measurements $\boldsymbol{y} = \boldsymbol{\Phi h}_0$, we match a subspace using a variation of (1):

$$\hat{\theta} = \arg\min_{\theta\in\Theta} \min_{\boldsymbol{h}\in\mathcal{S}_\theta} \|\boldsymbol{y} - \boldsymbol{\Phi h}\|_2^2 = \arg\min_{\theta\in\Theta} \|\boldsymbol{y} - \tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2, \tag{2}$$

where $\tilde{\boldsymbol{P}}_\theta$ is the orthogonal projection onto the range of $\boldsymbol{\Phi P}_\theta$, the K dimensional subspace of $\mathbb{R}^M$ consisting of measurements induced by signals in $\mathcal{S}_\theta$. We would like the compressed estimate produced by (2) to be close to the standard estimate in (1). We will judge the difference between these two estimates based on how well the subspaces they return can approximate the original signal $\boldsymbol{h}_0$. Our main result bounds the difference in the relative approximation errors $\hat{E}^2 - \bar{E}^2$, where

$$\bar{E}^2 = \frac{\|\boldsymbol{h}_0 - \boldsymbol{P}_{\bar{\theta}}\boldsymbol{h}_0\|_2^2}{\|\boldsymbol{h}_0\|_2^2} = 1 - \frac{\|\boldsymbol{P}_{\bar{\theta}}\boldsymbol{h}_0\|_2^2}{\|\boldsymbol{h}_0\|_2^2}, \quad \text{and}$$
$$\hat{E}^2 = \frac{\|\boldsymbol{h}_0 - \boldsymbol{P}_{\hat{\theta}}\boldsymbol{h}_0\|_2^2}{\|\boldsymbol{h}_0\|_2^2} = 1 - \frac{\|\boldsymbol{P}_{\hat{\theta}}\boldsymbol{h}_0\|_2^2}{\|\boldsymbol{h}_0\|_2^2}.$$

Since all of the terms above scale with the size of $\boldsymbol{h}_0$, we will assume (without loss of generality) from this point forward that $\|\boldsymbol{h}_0\|_2 = 1$, and derive a bound for the gap

$$\hat{E}^2 - \bar{E}^2 = \|\boldsymbol{P}_{\bar{\theta}}\boldsymbol{h}_0\|_2^2 - \|\boldsymbol{P}_{\hat{\theta}}\boldsymbol{h}_0\|_2^2. \tag{3}$$

Notice that the difference above must be positive, as $\bar{\theta}$ as given by (1) is the index for the optimal subspace. Note also that we are making no assumptions about whether or not $\boldsymbol{h}_0$ is in or even close to one of the $\mathcal{S}_\theta$. Bounding (3) also does not give us any immediate guarantees on how close the indices $\bar{\theta}$ and $\tilde{\theta}$ are to one another; we are only judging the effectiveness of the compressed program by its ability to produce a subspace which can estimate $\boldsymbol{h}_0$ almost as well as the optimal one. Given some assumptions about the collection of subspaces (say, that $\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|$ is small only when $\|\theta_1 - \theta_2\|_2$ is small), it may be possible to infer a bound on the size of $\hat{\theta} - \tilde{\theta}$ from a bound on (3), but we do not explore this here.

Our main results, stated as Theorems 1–3 in Section 2, provide a bound on this approximation gap whose most significant term has the form

$$\hat{E}^2 - \bar{E}^2 \;\lesssim\; \sqrt{\frac{K(\Delta + \log K)}{M}}, \tag{4}$$

where Δ is a quantity that captures the geometrical complexity of the collection of subspaces $\{\mathcal{S}_\theta\}$. As described in detail in Section 2, it is related to the covering number of $\{\mathcal{S}_\theta\}$ under the standard metric for subspaces, the operator norm of the difference of projectors $d(\theta_1, \theta_2) = \|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|$. In all of our motivating applications, we will have Δ = small constant + $\log(K)$. This means that we can control the approximation gap in (3) by making M slightly larger than K (and potentially significantly smaller than N). The result (4) holds "with high probability" for an arbitrary (but fixed) $\boldsymbol{h}_0$, where the randomness comes from the matrix $\boldsymbol{\Phi}$.

As this requires solving a least-squares problem of dimension K, it will only be practical to search over all of Θ when D (the dimension of the $\theta \in \Theta$) is small (1 or 2). This is indeed the case in the applications we discuss in the next section.

1.1 Motivating applications

1.1.1 Pulse detection from compressed samples

Perhaps the most fundamental example of a subspace matching problem is estimating the time shift (or "time of arrival") of a known pulse signal. In this case, the family $\mathcal{S}_\theta$ consists of one dimensional subspaces corresponding to scalings of a shifted template function $f_0(t)$; $\mathcal{S}_\theta = \{a\, f_0(t-\theta) \ : \ a \in \mathbb{R}\}$. Treatment of this problem with observations made through a random matrix $\mathbf{\Phi}$ was first given in [15,16], where it was called the *smashed filter*. An analysis of this of problem with a $\mathbf{\Phi}$ corresponding to measurements made in the frequency domain can be found in [20].

Figure 1 illustrates a variation on this problem for which we give a thorough analysis in Section 3.2. Here we are trying to find the shift and modulation frequency of a pulse of fixed shape that most closely matches the observed signal. The family of subspaces is indexed by the multi-parameter $\theta = (\tau, \omega)$, and consists of finely sampled[1] Gabor functions at different shifts, frequencies, and phases. We define

$$g(t; \tau, \omega, a, b) = e^{-(t-\tau)^2/\sigma^2} \left(a \cos(\omega t) + b \sin(\omega t)\right) \tag{5}$$

and $g[n; \tau, \omega, a, b] = g(t_1 + n\, d; \tau, \omega, a, b)$ for some sample spacing d and $n = 1, \ldots, N$; the vector $g[n; \cdot]$ contains equally spaced samples of $g(t; \cdot)$ on the interval $[t_1, t_2]$ with $t_2 = t_1 + Nd$. We can now write

$$\mathcal{S}_\theta = \{g[n; \tau, \omega, a, b] \ a, b \in \mathbb{R}, \ (\tau, \omega) = \theta \in \Theta\}. \tag{6}$$

Each subspace has a dimension of $K = 2$, and Θ is some subset of the time-frequency plane; here we will fix the specific region as $\Theta = [-0.25, 0.25] \times [2\pi \cdot 50, 2\pi \cdot 250]$. The example in Figure 1 uses a $\boldsymbol{h}_0$ consisting of (again finely spaced) samples of the signal $h(t) = (1 + \cos(\pi t/\beta_0)) \cdot \cos(2\pi \cdot 5t/\beta_0 + \pi/3) \cdot 1_{|t| \le \beta_0}$ where $\beta_0 = 5/128$. Notice that $\boldsymbol{h}_0$ is not actually included in any of the subspaces in $\mathcal{S}_\theta$; but the closest match does indeed comes at the expected values of $\hat{\tau} = 0$ and $\hat{\omega} = 2\pi \cdot 128$.

We see from the figure that gap between the uncompressed energy surface $\|\boldsymbol{h}_0 - \boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2$ and the compressed energy surface $\|\boldsymbol{y} - \tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2$ gets tighter as M increases. This is mathematically characterized by our main theorems along with the computations in Section 3.2.

1.1.2 Fast forward modeling for source localization

Another application of compressed subspace matching is accelerating modeling algorithms for source localization. For example, [32] shows how a *matched-field processing* (MFP) algorithm, working directly in the compressed domain, can be used to locate an underwater acoustic target in range and depth based on recorded frequency-domain acoustic measurements at a series of receivers at known locations. Although the spectral content of the source is generally unknown, there exists a model for the frequency response, the *Green's function* $\boldsymbol{G}_\theta(\omega) \in \mathbb{C}^N$, between any hypothetical source location $\theta = (r, z)$ (containing range and depth) and receivers $n = 1, \ldots, N$. If a source at frequency ω and complex amplitude $\beta(\omega) \in \mathbb{C}$ is active at location θ, the response at the N receivers is modeled as

$$\boldsymbol{h}(\omega) = \beta(\omega) \boldsymbol{G}_\theta(\omega) + \text{noise}.$$

[1] As we will discuss further in Section 3, the number of samples N plays no role in our analysis. We will simply assume that is is large enough so that the relationships between the $g(t; \cdot)$ are preserved.

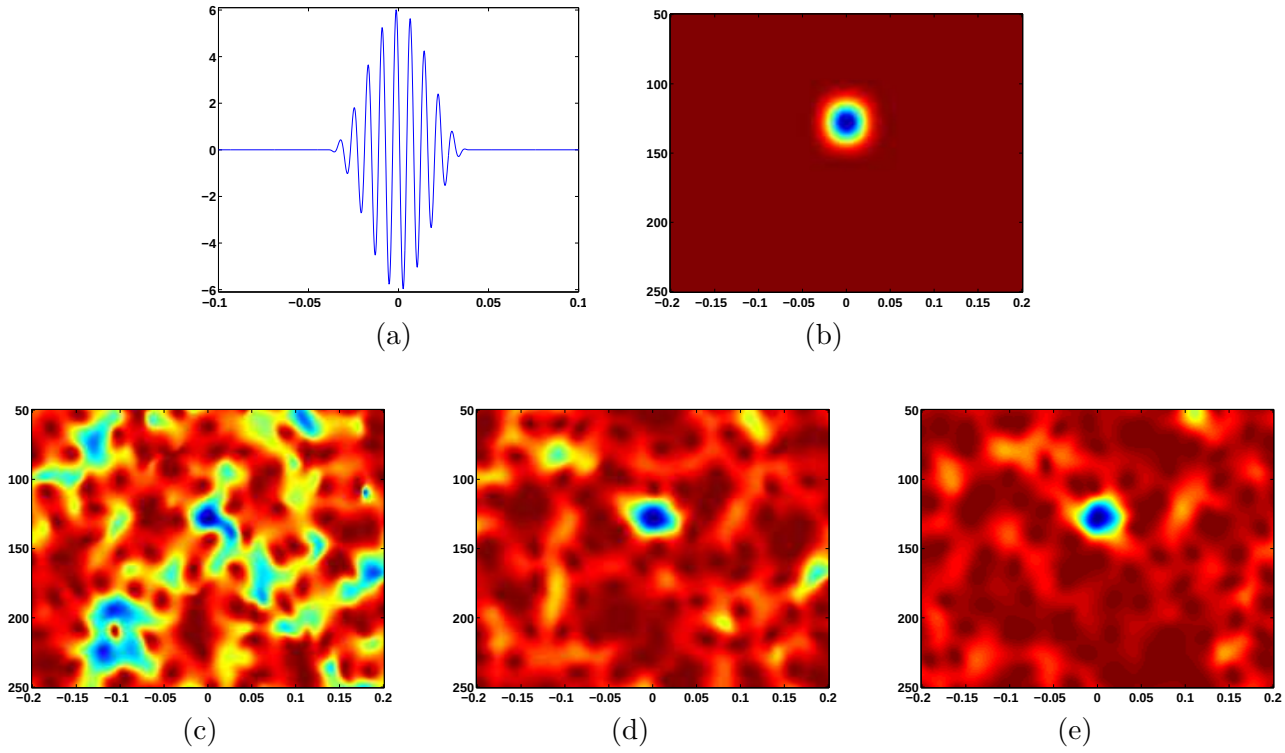

Figure 1: *(a) The original signal $\boldsymbol{h}_0$ is a raised cosine window on the interval $|t| \leq 5/128$ and is modulated to frequency $\omega_0 = 2\pi \cdot 128$. (b) The energy difference image $\|\boldsymbol{h}_0 - \boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2$. Each point in this image is the difference between the energy in $\boldsymbol{h}_0$ and the energy of its orthogonal projection onto a two dimensional subspace corresponding to a particular shift and modulation frequency, as in* (5)–(6)*. (c) One realization of the compressed energy difference image $\|\boldsymbol{y} - \tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2$ formed from $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{h}_0$ with $M = 10$; (d) a realization with $M = 20$; (e) $M = 30$. The minima in (c),(d), and (e) are all in exactly the same location as in (b), but clearly gets more distinct as M increases.*

Given observations at one or more frequencies[2], we can estimate the source location by solving (1) where the (real-valued, two-dimensional) subspaces are given by the column space of the $2N \times 2$ matrices

$$\boldsymbol{G}'_\theta = \begin{bmatrix} \mathrm{Re}\,(\boldsymbol{G}_\theta) & -\mathrm{Im}\,(\boldsymbol{G}_\theta) \\ \mathrm{Im}\,(\boldsymbol{G}_\theta) & \mathrm{Re}\,(\boldsymbol{G}_\theta) \end{bmatrix}$$

Typically, the responses $\boldsymbol{G}_\theta$ are only known implicitly, coming from higher-level knowledge about the geometry and materials in the channel. They can be acquired through simulating a source at each receiver location, and backpropagating the response over the locations of interest; if there are many receivers and the simulations are done across many frequencies, these simulations can be computationally expensive. In [32], it was demonstrated that the matched field processing could be done effectively from a compressed version of the $\boldsymbol{G}_\theta$. Sources with randomly generated amplitudes were simulated at all receiver locations simultaneously, resulting in a series of indirect observations $\boldsymbol{\phi}_1^{\mathrm{T}} \boldsymbol{G}'_\theta, \ldots, \boldsymbol{\phi}_M^{\mathrm{T}} \boldsymbol{G}'_\theta$ of the Green's function. These computations can be stacked on top of one another to form $\boldsymbol{\Phi}\boldsymbol{G}'_\theta$ for all θ — then a set of given observations can be encoded as $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{h}$ and the compressed version (2) of the problem is solved.

[2]If the source signature is known across frequency, the responses $\boldsymbol{G}_\theta(\omega)$ are added coherently into a vector of length N; if the signature is unknown, the responses are concatenated into a longer vector.

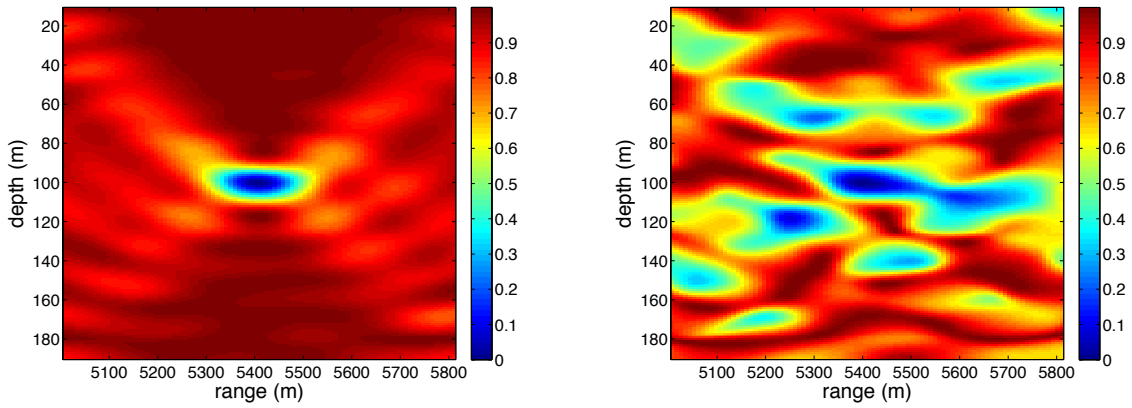

(a) (b)

Figure 2: *A subspace matching problem for an underwater acoustic channel. (a) The energy difference* $\|\boldsymbol{h}_0 - \boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2$ *for source located at* $\theta = (r, z) = (5405\text{m}, 100\text{m})$*. This image was formed using full knowledge of the (real-valued) Green's function* $\boldsymbol{G}'_\theta(\omega)$ *at single frequency. (b) The difference in the compressed domain formed using* $\boldsymbol{\Phi}\boldsymbol{G}'_\theta$ *with* $M = 6$*.*

Figure 2 shows the classical and compressive error surfaces for a shallow-water channel as a function of range and depth using $M = 6$ compressive measurements. Here, the Green's function is modeled as a Pekeris waveguide [26] at a frequency of 150 Hz with $N = 37$ uniformly spaced receivers along a vertical array between 10 meters and 190 meters in depth. Note that although the compressive error surface is a rough approximation to the true error surface, they are close enough to share a characteristically similar minimizer.

1.2 Linear encoding of a collection of subspaces

A key part of analyzing the direct-versus-compressed performance gap in (3) is showing that under appropriate conditions on the number of rows, the random $M \times N$ matrix $\boldsymbol{\Phi}$ provides a *stable embedding* for the collection of subspaces $\{\mathcal{S}_\theta\}$. Theorem 2 below says that if $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ are both in (possibly different) subspaces in $\{\mathcal{S}_\theta\}$, say $\boldsymbol{x}_1 \in \mathcal{S}_{\theta_1}$, $\boldsymbol{x}_2 \in \mathcal{S}_{\theta_2}$, then with high probability, the linear mapping $\boldsymbol{\Phi}$ preserves their distance. That is, there exists a $\delta > 0$ such that

$$(1-\delta)\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2^2 \;\le\; \|\boldsymbol{\Phi}(\boldsymbol{x}_1 - \boldsymbol{x}_2)\|_2^2 \;\le\; (1+\delta)\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2^2 \quad \boldsymbol{x}_1 \in \mathcal{S}_{\theta_1},\ \boldsymbol{x}_2 \in \mathcal{S}_{\theta_2},\ \text{for all } \theta_1, \theta_2 \in \Theta. \tag{7}$$

In fact, Theorem 2 shows that we can take

$$\delta \;\lesssim\; \sqrt{\frac{K(\Delta + \log K)}{M}}, \tag{8}$$

where Δ is the same measure of geometric complexity of $\{\mathcal{S}_\theta\}$ as in the subspace matching bound in (4). This means that the embedding is uniformly stable for all θ_1, θ_2 when the number of rows in $\boldsymbol{\Phi}$ is just a little larger than the dimension K of each subspace.

Random projections are known to be efficient stable linear embeddings in sparse signal models, manifolds, etc., so it instructive to compare the content of (7) and (8) to previous results in the

literature. For an arbitrary $\mathcal{Q} \subset \mathbb{R}^N$, let $\delta(\mathcal{Q})$ be defined as

$$\delta(\mathcal{Q}) = \max_{\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{Q}} \frac{\left| \|\boldsymbol{\Phi}(\boldsymbol{x}_1 - \boldsymbol{x}_2)\|_2^2 - \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2^2 \right|}{\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2^2}.$$

When $\mathcal{Q}$ is a finite set of size $|\mathcal{Q}| = Q$, then a now classical result knows as the *Johnson-Lindenstrauss* Lemma [13, 27] tells us that we can take

$$\delta(\mathcal{Q}) \lesssim \sqrt{\frac{2 \log Q}{M}},$$

meaning that $\boldsymbol{\Phi}$ preserves distances when $M \gtrsim \log Q$.

An analog to the Johnson-Lindenstrauss Lemma has been developed for arbitrary sets $\mathcal{Q}$ [17, 28, 35] (i.e. sets with infinite cardinality). In these works, the $\log Q$ above is (approximately speaking) replaced with the *Gaussian width*. These results are closely related to our Theorem 2, with the main difference being that our bounds explicitly rely on the geometry of the set of subspaces, rather than a set of points.

When $\mathcal{Q}$ is a linear subspace of dimension K, then $\delta(\mathcal{Q})$ is directly related to the singular values of $\boldsymbol{\Phi}$; in particular, it is the maximum deviation of the singular spectrum from 1. Random matrices of this kind are well conditioned when they are overdetermined by a constant factor [14, 33, 42]; we can take

$$\delta(\mathcal{Q}) \lesssim \sqrt{\frac{K}{M}}.$$

A straightforward extension [2, 11, 42] of this result yields a uniform embedding guarantee when $\mathcal{Q}$ is all vectors in a finite collection of subspaces $\{\mathcal{S}_\theta \ : \ \theta \in \Theta\}$. In this case,

$$\delta(\mathcal{Q}) \lesssim \sqrt{\frac{2K + 2 \log |\Theta|}{M}}, \tag{9}$$

where $|\Theta|$ is the number of subspaces. In the literature on compressive sensing, where the goal is to recover a sparse vector from linear measurements through $\boldsymbol{\Phi}$, stable recovery for all K-sparse vectors can be guaranteed [8, 10] by considering $\mathcal{Q}$ to be the set of all vectors in $\mathbb{R}^N$ with at most K active components, a union of subspaces of size $|\Theta| = \binom{N}{K} \lesssim e^{K \log(N/K)}$. Sparse vectors, then, can be encoded using $M \gtrsim K \log(N/K)$ random measurements.

In this paper, we are interested in collections of subspaces indexed by θ on the continuum, meaning that $|\Theta|$ is (uncountably) infinite. This requires a more nuanced treatment than simply applying (9). The bound in (9) is derived by treating each subspace individually, and so depends only on the number of subspaces to treat. We will more carefully take advantage of the relationships between subspaces that are very close to one another. The geometric constant Δ in (8) is, roughly speaking, a measure of how well the collection $\{\mathcal{S}_\theta\}$ can be approximated by a finite cover across many resolutions.

There is another important result in the sparse and low-rank recovery literature that explicitly handles a continuum of subspaces. In [36], a uniform embedding result was presented for the set of $N \times N$ matrices of rank at most K (this result was refined in [9]) passing through a random projection (in this case, $\boldsymbol{\Phi}$ would be $M \times N^2$). The argument essentially reduces to taking $\mathcal{Q}$ to be *all* (pairs of) K dimensional subspaces of $\mathbb{R}^N$, with

$$\delta(\mathcal{Q}) \lesssim \sqrt{\frac{NK}{M}}.$$

In this paper, we consider (a possibly very small) subset of these spaces, and so in general do not require M to be as large.

1.3 Related work

The compressed subspace matching problem, and its analysis put forth in this paper, is closely related to *compressed sensing*, where the goal is to recover a structured (i.e. sparse) vector from an underdetermined (and usually random) linear observations. As mentioned in Section 1.2, many of the results in this field hinge on a measurement matrix $\mathbf{\Phi}$ stably embedding a collection of subspaces [10, 11, 19]. Embedding results have been derived for different finite collections of subspaces of interest to different applications, including sparse signals [2], signals with certain types of "group sparsity" [6, 22], and arbitrary finite-size collections [5, 23]. In contrast to these works, this paper focusses on a union of subspaces indexed by a continuous parameter. We could of course sample θ on a grid and apply these known results, but the bounds acquired this way would depend on the size of the grid chosen (which is deeply unsatisfying in that it violates the spirit of the continuum), and there are invariably complications with values of θ that are "off the grid" (between grid points). Our approach, taking θ as a point in the continuum, is simply more direct. The recent work [38], which looks at the particular case of estimating the frequencies of sinusoids, works from this same perspective.

The recent paper [4] also discusses the general problem of recovering a signal which belongs to an arbitrary union of subspaces. This work essentially shows that if $\mathbf{\Phi}$ is an embedding over this space, then the signal can be reconstructed using a projected Landweber algorithm, where each iteration includes a solution of (1). This paper provides guarantees for the direct solution of (2) — on the continuum, (1) and (2) are comparable computationally — and provides conditions under which $\mathbf{\Phi}$ is likely to provide such an embedding.

As we discuss in Section 1.1, a scenario in which solving the subspace matching problem is of particular interest is in estimating the time-of-arrival of a pulse. Specialized algorithms for this particular problem have also been developed in [25] and [44]. More general algorithms for recovering signals from unions of subspaces can be found in [21, 24].

Other closely related recent results come from the area of manifold embeddings. In [3], it is shown that pairwise distances between points on a manifold are preserved through a random projection; they derive a sufficient condition on the dimension of this projection which involves the volume of the manifold, its ambient and intrinsic dimensions, and its curvature. This result was later refined in [12], which removed the dependence on the ambient dimension and used a more refined definition of curvature. Variations on these results for a variety of different projection operators can be found in [43]. In this work, we are explicitly concerned with collections of subspaces that are smoothly parameterized rather than a collection of vectors. One advantage of this is it allows for a more natural treatment of scale — we do not have any size constraints on our signals of interest. Similar to these manifold embedding results, our bounds do depend on quantifying the geometry of the set of interest. The main geometrical concept we use, the size of an finite approximation of the collection of subspaces to a certain precision, is different than the descriptors typically used to manifolds (curvature, volume, etc).

One of the keys to establishing that the gap in (3) is small is showing that $\mathbf{\Phi}$ preserves the distances between the subspaces we are considering; this is the content of Theorem 2 in Section 2. As

mentioned in Section 1.2, there are closely related results in the probability literature. In [28], a type of Johnson-Lindenstrauss lemma for arbitrary sets was put forth where the geometry of the set was measured using the "generic chaining complexity". The methods developed there were refined in [34, 35], and an explicit relationship between the dimension of the projection and the accuracy of the embedding was proven in [17]. In [7], similar embedding results are established for sparse random matrices. In this work, we explicitly quantify the geometrical complexity of the collection of subspaces in terms of the natural metric (operator norm of the distance between projectors) which allows us to calculate the bounds explicitly for several scenarios of immediate interest in signal processing (see Section 3).

Very recent and independent work [18] applies advanced chaining techniques directly to the problem of embedding a collection of subspaces through a random matrix. It is possible that the results in this work could remove the logarithmic factor in Theorem 2.

2 Main results

For a fixed $\boldsymbol{h}_0$, we wish to bound the difference in the approximation error $\hat{E}^2$ of the subspace matched from indirect measurements $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{h}_0$ and the approximation error $\bar{E}^2$ of the subspace matched from a direct observation. (The direct observation error can be thought of as an oracle error). Our analysis will assume that the matrix $\boldsymbol{\Phi}$ is Gaussian, with

$$\Phi[m,n] \;\sim\; \text{Normal}\left(1, M^{-1}\right), \tag{10}$$

and all of the entries being independent.

To do this, we will establish a *uniform bound* on the difference between the energy of the projection of $\boldsymbol{h}_0$ onto subspace $\mathcal{S}_\theta$ is close to the energy of the projection of $\boldsymbol{y}$ onto the range of $\boldsymbol{\Phi}\boldsymbol{P}_\theta$. We will show that

$$\sup_{\theta\in\Theta}\left| \|\tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| \;\leq\; \delta, \tag{11}$$

where δ depends on the number of rows M and the geometry of the collection of subspaces $\{\mathcal{S}_\theta\}$. It immediately follows that

$$\begin{aligned}
\hat{E}^2 - \bar{E}^2 &= \|\boldsymbol{P}_{\bar{\theta}}\boldsymbol{h}_0\|_2^2 - \|\boldsymbol{P}_{\hat{\theta}}\boldsymbol{h}_0\|_2^2 \\
&\leq \|\tilde{\boldsymbol{P}}_{\hat{\theta}}\boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_{\hat{\theta}}\boldsymbol{h}_0\|_2^2 + \|\boldsymbol{P}_{\bar{\theta}}\boldsymbol{h}_0\|_2^2 - \|\tilde{\boldsymbol{P}}_{\bar{\theta}}\boldsymbol{y}\|_2^2 \\
&\leq 2\delta,
\end{aligned}$$

where the first inequality follows from the fact that $\|\tilde{\boldsymbol{P}}_{\hat{\theta}}\boldsymbol{y}\|_2^2 > \|\tilde{\boldsymbol{P}}_{\bar{\theta}}\boldsymbol{y}\|_2^2$ by the definition of $\hat{\theta}$ in (2),

$$\hat{\theta} = \arg\min_{\theta\in\Theta} \|\boldsymbol{y} - \tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2 = \arg\max_{\theta\in\Theta} \|\tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2.$$

With $\boldsymbol{\Phi}$ random, we can interpret (11) as a bound on the supremum of a *random process* indexed by $\theta \in \Theta$. Our upper bound δ will not hold in an absolute sense, but rather with an appropriately high probability. The supremum in (11) follows from bounding the suprema of two related random processes:

$$\text{(C1)} \qquad \sup_{\theta\in\Theta} \|\boldsymbol{P}_\theta - \boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_\theta\| \leq \delta_1,$$

$$\text{(C2)} \qquad \sup_{\theta\in\Theta} \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_\theta^{\perp}\boldsymbol{h}_0\|_2 \leq \delta_2.$$

The first bound (C1) tells us that all of the projection operators $\boldsymbol{P}_\theta$ retain their structure when viewed through the matrix $\boldsymbol{\Phi}$. The norm in (C1) is the standard operator norm (largest magnitude eigenvalue in this case). The second bound (C2) tells us that $\boldsymbol{P}_\theta^\perp \boldsymbol{h}_0$ stays approximately orthogonal to the subspace $\mathcal{S}_\theta$ after being passed through $\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}$.

Our first main result, which we prove in Section 4.1 relates (C1) and (C2) to our main goal (11).

Theorem 1 *For a matrix $\boldsymbol{\Phi}$ and collection of subspaces $\{\mathcal{S}_\theta,\ \theta \in \Theta\}$ as described above, if (C1) and (C2) hold with $\delta_1, \delta_2 < 1$, then*

$$\sup_{\theta\in\Theta} \left| \|\tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| \ \leq\ \delta := \frac{3\delta_1 + 2\delta_2 + (\delta_1+\delta_2)^2}{1-\delta_1}.$$

The upper bounds in (C1) and (C2) depend on the *geometrical structure* of the collection of subspaces $\{\mathcal{S}_\theta\}$. This structure is quantified through how well $\{\mathcal{S}_\theta\}$ can be approximated at different resolutions using a finite subset of the subspaces. This approximation is taken using the natural distance between two subspaces, the operator norm of the difference of their projectors:

$$d(\mathcal{S}_{\theta_1}, \mathcal{S}_{\theta_2}) \triangleq \|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|.$$

Notice that it is always true that $0 \leq d(\mathcal{S}_{\theta_1}, \mathcal{S}_{\theta_2}) \leq 1$. A finite subset $\{\mathcal{T}_q := \mathcal{S}_{\theta_q},\ \ \{\theta_1, \ldots, \theta_Q\} \subset \Theta\}$ is an ϵ-approximation to $\{\mathcal{S}_\theta\}$ if everything in $\{\mathcal{S}_\theta\}$ is within distance ϵ of a corresponding element in $\{\mathcal{T}_q\}$,

$$\sup_{\theta\in\Theta} \min_{1\leq q\leq Q} d(\mathcal{S}_\theta, \mathcal{T}_q) \ \leq\ \epsilon.$$

The covering number of $\{\mathcal{S}_\theta\}$ at resolution ϵ it the size of the smallest such ϵ-approximation

$$N(\{\mathcal{S}_\theta\}, \epsilon) = \min \left\{ \mathrm{card}(\{\mathcal{T}_q\}) \ :\ \sup_{\theta\in\Theta} \min_q d(\mathcal{S}_\theta, \mathcal{T}_q) \ \leq\ \epsilon \right\}.$$

Of course, as $\epsilon \to 0$, the covering number $N(\{\mathcal{S}_\theta\}, \epsilon)$ increases. Our assumption about the collection of subspaces $\{\mathcal{S}_\theta\}$ is that this covering number grows at most polynomially as ϵ decreases. That is, there exist constants N_0 and α such that

$$N(\{\mathcal{S}_\theta\}, \epsilon) \ \leq\ N_0\, \epsilon^{-\alpha}, \quad \text{for all} \quad 0 \leq \epsilon \leq 1. \tag{12}$$

We call such collections of subspaces *geometrically regular*. As it turns out, this assumption is met for many and perhaps even most cases of practical interest. In Section 3, we compute bounds on (N_0, α) in a few cases where Θ is a subset of the real line that indexes the position (shift) of K orthogonal basis functions. In each of these cases, we see that $N_0 \leq \mathrm{poly}(K) \cdot |\Theta|$ and $\alpha = 1$ or 2. In general, we will treat α like a small constant, with N_0 capturing the most important parts of the geometrical complexity of the set $\{\mathcal{S}_\theta\}$.

In the bounds below for (C1) and (C2), (N_0, α) are combined into a single constant Δ which has the form

$$\Delta = \log(8^\alpha N_0^2) + 2 \ \approx\ 2.08\,\alpha + 2\log N_0 + 2. \tag{13}$$

Roughly speaking, this quantity approximates (the logarithm of) the number of subspaces over which the processes $\|\boldsymbol{P}_\theta - \boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_\theta\|$ and $\|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_\theta^\perp \boldsymbol{h}_0\|_2$ are significantly different.

Our next main result relates gives us a means to control the upper bound δ_1 in (C1) in terms of these geometrical descriptors of $\{\mathcal{S}_\theta\}$. In both Theorems 2 and 3, we assume that $M \geq K$, which is a necessary condition to stably embed even a single subspace of dimension K.

Theorem 2 *Let $\mathbf{\Phi}$ be an $M \times N$ matrix with independent Gaussian entries as in* (10)*, and let $\{\mathcal{S}_\theta\}$ be a collection of subspaces that are geometrically regular in the sense of* (12) *with Δ defined as in* (13)*. Then there exist constants C_1 and C_2 such that for all $t \geq 0$*

$$\mathrm{P}\left(\sup_{\theta\in\Theta} \|\boldsymbol{P}_\theta - \boldsymbol{P}_\theta \mathbf{\Phi}^{\mathrm{T}} \mathbf{\Phi} \boldsymbol{P}_\theta\| > C_1\, \gamma(t)\right) \;\leq\; C_2 K\, e^{-t+\Delta}, \tag{14}$$

where

$$\gamma(t) = \sqrt{\frac{Kt}{M}} + \frac{Kt}{M} + \frac{t}{\sqrt{M}} + \frac{t^2}{M} + \sqrt{\frac{K}{M}} + \frac{\alpha^2}{M}. \tag{15}$$

To make the probability bound in (14) meaningful, we will need to take $t \gtrsim \Delta + \log K$. It follows that, with high probability, we have the following upper bound on δ_1:

$$\delta_1 \;\lesssim\; \left(\frac{\max\left(K(\Delta + \log K), (\Delta + \log K)^2\right)}{M}\right)^{1/2}.$$

This bound will be meaningful when we have:

$$M \;\gtrsim\; \max\left(K(\Delta + \log K), (\Delta + \log K)^2\right).$$

As we will see in the examples in Section 3, it is typical that Δ scales like $\log K$ as well.

This δ_1 term also gives an immediate bound on the embedding all signals from these subspaces with respect to the random projection $\mathbf{\Phi}$. To wit, for any δ_1 satisfying (C1), we have:

$$\sup_{\theta\in\Theta}\sup_{f\in\mathcal{S}_\theta}\left|\frac{\|\mathbf{\Phi} f\|^2}{\|f\|^2} - 1\right| = \sup_{\theta\in\Theta}\sup_{\|f\|=1} \|f^T \boldsymbol{P}_\theta f - f^T \boldsymbol{P}_\theta \mathbf{\Phi}^{\mathrm{T}} \mathbf{\Phi} \boldsymbol{P}_\theta f\| = \sup_{\theta\in\Theta} \|\boldsymbol{P}_\theta - \boldsymbol{P}_\theta \mathbf{\Phi}^{\mathrm{T}} \mathbf{\Phi} \boldsymbol{P}_\theta\| \leq \delta_1, \tag{16}$$

mirroring the classic expression for the restricted isometry property of sparse vectors.

Our final main result is a tail bound of a similar form for the upper bound δ_2 in (C2).

Theorem 3 *With $\mathbf{\Phi}$ and $\{\mathcal{S}_\theta\}$ as in Theorem 2, there exist constants C_1 and C_2 such that for any fixed $\boldsymbol{h}_0 \in \mathbb{R}^N$ with $\|\boldsymbol{h}_0\|_2 = 1$,*

$$\mathrm{P}\left(\sup_{\theta\in\Theta} \|\boldsymbol{P}_\theta \mathbf{\Phi}^{\mathrm{T}} \mathbf{\Phi} \boldsymbol{P}_\theta^{\perp} \boldsymbol{h}_0\| > C_1\, \gamma(t)\right) \;\leq\; C_2 K e^{-t+\Delta}$$

where

$$\gamma(t) = \sqrt{\frac{tK}{M}} + \frac{t\sqrt{K}}{M} + \sqrt{\frac{K}{M}}.$$

As with Theorem 2, we can make the probability above small by taking $t \gtrsim \Delta + \log K$, meaning that $\sup_{\theta\in\Theta} \|\boldsymbol{P}_\theta \mathbf{\Phi}^{\mathrm{T}} \mathbf{\Phi} \boldsymbol{P}_\theta^{\perp} \boldsymbol{h}_0\| \leq \delta < 1$ for $M \gtrsim \delta^{-2} K(\Delta + \log K)$.

3 Collections of subspaces with geometric regularity

In this section, we will discuss some specific examples of a parameterized collection of subspaces $\{\mathcal{S}_\theta\}$ that are geometrically regular in that there exist constants N_0, α such that

$$N(\{\mathcal{S}_\theta\}, \epsilon, d) \leq N_0 \epsilon^{-\alpha}, \quad d(\mathcal{S}_{\theta_1}, \mathcal{S}_{\theta_2}) = \|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\| \quad \forall \epsilon \leq 1.$$

In each of the examples, mapping from $\theta \in \mathbb{R}^D$ to the K-dimensional subspace $\mathcal{S}_\theta$ is Holder-continuous. For $D = 1$, this means that

$$\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\| \leq \beta|\theta_1 - \theta_2|^\rho, \tag{17}$$

for some constants β and ρ. This allows us to directly tie the covering numbers of the set of subspaces $\{\mathcal{S}_\theta\}$ to the covering numbers of the parameter class Θ under the standard Euclidean metic as

$$\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\| \leq \beta|\theta_1 - \theta_2|^\rho \;\Rightarrow\; N(\{\mathcal{S}_\theta\}, \epsilon, d) \leq N(\Theta, (\epsilon/\beta)^{1/\rho}, |\cdot|). \tag{18}$$

The sets Θ we consider are simple: in two of the three examples below, they are just intervals on the real line. In this case, we have the easy bound

$$N([a,b], \epsilon, |\cdot|) \leq \frac{|b-a|}{2\epsilon} + 1 \leq \frac{|b-a|}{\epsilon}, \quad \forall \epsilon \leq (1/\beta)^{1/\rho},$$

provided that $|b-a| \geq (\epsilon/\rho)^{1/\rho}$. So if, for example, we establish (17) and $\Theta = [a,b]$, we can take $\alpha = 1/\rho$ and $N_0 = \beta^{1/\rho}|b-a|$.

All of our examples involve collections of subspaces index by a shift parameter or (equivalently) a modulation parameter. The general idea is that the signal lies in the span (or can be closely approximated by) a superposition of K basis functions that are localized in time; the parameter θ indicates the location of the basis functions. More precisely, for the shift parameter spaces, we are interested in signals that can be written

$$h(t) = \sum_{k=1}^{K} h[k] v_k(t-\theta), \quad \text{for some } \boldsymbol{h} \in \mathbb{R}^K \text{ and } \theta \in [t_1, t_2],$$

where $v_1(t), \ldots, v_K(t) \in L_2(\mathbb{R})$ are linearly independent, continuous-time signals. The collection of all such signals form a manifold in $L_2(\mathbb{R})$. Our model will be that we observe this signal (through Φ) after it has been projected onto a fixed N dimensional ambient space $\mathcal{H}$. The subspaces $\mathcal{S}_\theta \subset \mathcal{H}$ are then the span of the $v_k(t-\theta)$ projected onto $\mathcal{H}$. Note that we are **not** discretizing the parameter θ; we are simply articulating the signal model as a subset of $\mathbb{R}^N$ rather than $L_2(\mathbb{R})$. The exact nature of this discretization is not that critical for our purposes here; we simply introduce it to avoid problematic topological issues that are ultimately irrelevant to our understanding. The ambient dimension N does not appear anywhere in our results, and so this discretization can be arbitrarily fine-grained. We will perform our calculations below based on the $L_2(\mathbb{R})$ distance between continuous-time signals, and then use the simple fact that $\|\boldsymbol{P}_{\mathcal{H}}[u(t) - w(t)]\|_2 \leq \|u(t) - w(t)\|_{L_2(\mathbb{R})}$.

In each of our examples below, the $\{v_k(t-\theta)\}$ are a well-defined orthobasis for each subspace in the collection (before it is projected into $\mathcal{H}$). We bound the distance between the subspaces parameterized by θ_1 and θ_2 by looking at the distances between these basis functions at different shifts. If we set $u_1(t) = v_1(t-\theta_1), \ldots, u_K(t) = v_K(t-\theta_1)$ and $w_1(t) = v_1(t-\theta_2), \ldots, w_K(t) = v_K(t-\theta_2)$, then we have

$$\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\| \leq 2 \sup_{\|x(t)\|_2 = 1} \left(\sum_{k=1}^{K} |\langle x(t), u_k(t) - w_k(t) \rangle|^2 \right)^{1/2} \leq 2 \left(\sum_{k=1}^{K} \|u_k(t) - w_k(t)\|_2^2 \right)^{1/2}. \tag{19}$$

The examples here are meant to be illustrative more than precise. The geometrical constant Δ in Theorems 2 and 3 scale like $\Delta \sim \alpha + \log N_0$, and so our main concern will be establishing α as a small constant, and N_0 a polynomial of K of modest order.

3.1 Pulse at an unknown delay

In our first example, we consider the problem of estimating the shift (or "time of arrival") of a single pulse with a known shape. In this case, solving the direct observation problem (1) is equivalent to the classical *matched filter*; the indirect observation problem in (2) has been called the *smashed filter* in the literature [15, 16] and has been previously analyzed for specialized $\mathbf{\Phi}$ [20].

We start by considering a Gaussian pulse,

$$v(t) = \pi^{-1/4}\sigma^{-1/2}e^{-t^2/2\sigma^2},$$

where σ is a fixed width parameter and the normalization factor in front ensures $\|v(t)\|_2 = 1$. The collection of $K = 1$ dimensional subspaces we want to test are generated by the manifold of unit vectors $\{v(t-\theta),\ \theta \in [0,T]\}$ projected onto the discretized space $\mathcal{H}$. We have

$$\begin{aligned}\|v(t-\theta_1) - v(t-\theta_2)\|_2^2 &= \|v(t-\theta_1)\|_2^2 + \|v(t-\theta_2)\|_2^2 - 2\langle v(t-\theta_1), v(t-\theta_2)\rangle \\ &= 2\left(1 - e^{-(\theta_2-\theta_1)^2/4\sigma^2}\right) \\ &\leq \frac{|\theta_2-\theta_1|^2}{2\sigma^2},\end{aligned} \tag{20}$$

and so applying (19),

$$\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\| \leq \frac{\sqrt{2}}{\sigma}|\theta_2 - \theta_1|,$$

which combined with (18) means

$$N(\{\mathcal{S}_\theta\}, \epsilon, d) \ \leq \ \frac{\sqrt{2}\,T}{\sigma}\epsilon^{-1}.$$

In this case, we can use $\alpha = 1$ and $N_0 = \sqrt{2}T/\sigma$, and take

$$\Delta \leq \log(T/\sigma) + 4.78,$$

as the geometrical constant in Theorems 2 and 3.

As the calculations in (20) suggest, the distance between two close shifts of the same function is related to the flatness of its autocorrelation function near zero. For smooth $v(t)$, the following lemma, proved in the appendix, quantifies this behavior in a convenient manner.

Lemma 1 *Let $v(t)$ be a unit-energy signal, $\|v(t)\|_{L_2} = 1$, with Fourier transform $\hat{v}(\xi)$. If $v(t)$ has bounded Sobolev norm,*

$$\left(\frac{1}{2\pi}\int_{-\infty}^{\infty}\xi^2|\hat{v}(\xi)|^2\ d\xi\right)^{1/2} \ \leq \ L,$$

then

$$\|v(t-\theta_1) - v(t-\theta_2)\|_2 \ \leq \ L|\theta_1 - \theta_2|.$$

So if the collection of subspaces $\{\mathcal{S}_\theta\}$ consists of all shifts $\{v(t-\theta),\ \theta \in [0,T]\}$ of a Sobolev differentiable function (projected onto $\mathcal{H}$), then

$$N(\{\mathcal{S}_\theta\}, \epsilon, d) \leq LT\,\epsilon^{-1} \quad \Rightarrow \quad \Delta \leq 2\log(LT) + 4.08. \tag{21}$$

If the pulse $v(t)$ is not smooth, we can still bound the distance between shifts using the total variation of $v(t)$. Roughly speaking, the total variation is the arc length of the curve that $v(t)$ traces out; technically, we define the total variation of a signal on the real line as

$$\|v(t)\|_{TV} = \lim_{a\to\infty} \sup_{\mathcal{P}} \left\{ \sum_{i=1}^{|\mathcal{P}|} |v(t_{i+1}) - v(t_i)|,\ \mathcal{P} = \{t_1, \ldots, t_{|\mathcal{P}|}\} \text{ is a partition of } [-a, a]. \right\}.$$

The following lemma relates the total variation the energy in the difference between shifts.

Lemma 2 *Let $v(t)$ be a signal with finite total variation. Then*

$$\|v(t-\theta_1) - v(t-\theta_2)\|_2 \ \leq\ \|v(t)\|_{TV}\, |\theta_1 - \theta_2|^{1/2}.$$

As a quick application, suppose that $v(t)$ is a square pulse of length σ,

$$v(t) = \begin{cases} 1/\sqrt{\sigma}, & -\sigma/2 \leq t \leq \sigma/2, \\ 0 & \text{otherwise.} \end{cases}$$

Then $\|v(t)\|_{TV} = 2/\sqrt{\sigma}$, and

$$N(\{\mathcal{S}_\theta\}, \epsilon, d) \leq \frac{4T}{\sigma}\epsilon^{-2} \quad \Rightarrow \quad \Delta \leq 2\log(T/\sigma) + 8.94. \tag{22}$$

In cases where $v(t)$ is not smooth, we do expect $\{\mathcal{S}_\theta\}$ to be richer collection of subspaces, a fact which is reflected in the powers of ϵ in the bounds (21) and (22). However, the ultimate difference this makes for Δ is very minor — it simply results in a greater additive constant.

3.2 Gabor pulse, unknown shift and unknown modulation frequency

In our next example, we consider the problem of jointly estimating the shift and modulation frequency of a Gabor pulse:

$$g(t, \omega; a, b) = e^{-t^2/2\sigma^2} \left(a\cos(\omega t) + b\sin(\omega t)\right).$$

Each subspace in the collection has $K = 2$, and $\mathcal{S}_\theta$, where $\theta = (\tau, \omega)$ is a multiparameter with $D = 2$, consists of the projection onto $\mathcal{H}$ of the span of $e^{-(t-\tau)^2/2\sigma^2}\cos(\omega(t-\tau))$ and $e^{-(t-\tau)^2/2\sigma^2}\sin(\omega(t-\tau))$. For the search space, we take $\Theta = [0, T] \times [\Omega_l, \Omega_u]$. Define

$$v_1(t; \tau, \omega) = C_1(\omega) e^{-(t-\tau)^2/2\sigma^2}\cos(\omega(t-\tau)), \quad C_1(\omega) = \sqrt{\frac{2}{\sigma\sqrt{\pi}(1 + e^{-\omega^2\sigma^2})}}$$

$$v_2(t; \tau, \omega) = C_2(\omega) e^{-(t-\tau)^2/2\sigma^2}\sin(\omega(t-\tau)), \quad C_2(\omega) = \sqrt{\frac{2}{\sigma\sqrt{\pi}(1 - e^{-\omega^2\sigma^2})}}$$

We will assume that the pulse has at least a mild amount of oscillation by restricting $\Omega_l\sigma > 1$. We will bound the distances in two parts, using

$$\|v_1(t; \tau_1, \omega_1) - v_1(t; \tau_2, \omega_2)\|_2 \leq \|v_1(t; \tau_1, \omega_1) - v_1(t; \tau_2, \omega_1)\|_2 + \|v_1(t; \tau_2, \omega_1) - v_1(t; \tau_2, \omega_2)\|_2,$$

and do a similar calculation for v_2.

For the time shift, we use Lemma 1 along with the following expressions for the Fourier transforms:

$$\hat{v}_1(\xi;\tau,\omega) = \frac{1}{1+e^{-\omega^2\sigma^2}}\left(e^{-(\xi-\omega_1)^2\sigma^2/2} + e^{-(\xi+\omega_1)^2\sigma^2/2}\right)e^{-j\tau\xi}$$
$$\hat{v}_2(\xi;\tau,\omega) = \frac{1}{1-e^{-\omega^2\sigma^2}}\left(e^{-(\xi-\omega_1)^2\sigma^2/2} - e^{-(\xi+\omega_1)^2\sigma^2/2}\right)e^{-j(\tau\xi+\pi/2)}.$$

This leads to the bound

$$\|v_1(t;\tau_1,\omega_1) - v_1(t;\tau_2,\omega_1)\|_2 \leq \sqrt{\left(\omega_1^2 + \frac{1}{\sigma^2}\right)}\cdot|\tau_1-\tau_2|,$$

and the same bound holds for $\|v_2(t;\tau_1,\omega_1) - v_2(t;\tau_2,\omega_1)\|_2$. For the frequency shift,

$$\begin{aligned}\|v_1(t;\tau_2,\omega_1) - v_1(t;\tau_2,\omega_2)\|_2 &= (2\pi)^{-1/2}\|\hat{v}_1(\xi;\tau_2,\omega_1) - \hat{v}_1(\xi;\tau_2,\omega_2)\|_2 \\ &\leq \max(C_1(\omega_1),C_1(\omega_2))\sigma\left\|e^{-(\xi-\omega_1)^2\sigma^2/2} - e^{-(\xi-\omega_2)^2\sigma^2/2}\right\|_2 \\ &\leq \sigma|\omega_1-\omega_2|,\end{aligned}$$

where the last step follows from a similar calculation to the one at the beginning of Section 3.1. The calculation for the frequency shift of the v_2 is only slightly different,

$$\begin{aligned}\|v_2(t;\tau_2,\omega_1) - v_2(t;\tau_2,\omega_2)\|_2 &\leq \max(C_2(\omega_1),C_2(\omega_2))\,\sigma\left\|e^{-(\xi-\omega_1)^2\sigma^2/2} - e^{-(\xi-\omega_2)^2\sigma^2/2}\right\|_2 \\ &\leq 1.582\,\sigma|\omega_1-\omega_2|,\end{aligned}$$

where we have used the fact that $\Omega_l\sigma \geq 1$ to bound the maximum in the last inequality.

We can now combine these inequalities to bound the distance between the projectors. With $\theta_1 = (\tau_1,\omega_1)$ and $\theta_2 = (\tau_2,\omega_2)$, we have

$$\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\| \leq 2\sqrt{2}\left(\left(\Omega_u^2 + \frac{1}{\sigma^2}\right)^{1/2}|\tau_1-\tau_2| + 1.582\sigma|\omega_1-\omega_2|\right)$$

With $\Omega = \Omega_u - \Omega_l$ as the length of the search interval in frequency and T length of the search interval in time, we can then bound the covering numbers as

$$N(\{\mathcal{S}_\theta),\epsilon,d) \;\leq\; \frac{9\sqrt{\sigma^2\Omega_u^2+1}\,T\,\Omega}{\epsilon^2} \quad\Rightarrow\quad \Delta \leq 2\log((\Omega_u^2\sigma^2+1)^{1/2}T\Omega) + 8.36.$$

3.3 Multifrequency pulse, unknown shift

Consider the problem of trying to find a pulse (again at a location in $[0,T]$) which is localized in time and exists over a known frequency band. We might model such a signal as a superposition of modulations of a fixed envelope function $g(t)$:

$$h(t) = \sum_{k=1}^{K}\beta_k\,v_k(t), \quad \text{with } v_k(t) = g(t)\cos(\omega_k t).$$

With a careful choice of $g(t)$ and frequencies ω_k, the $v_k(t)$ are orthogonal; this is known as the *lapped orthogonal transform* [31]. To illustrate, we consider the particular case of

$$g(t) = \sqrt{\frac{2}{\sigma}} \begin{cases} 0: & t < -\sigma/4 \\ \sin\left(\frac{\pi}{4}\left(1+\frac{4t}{\sigma}\right)\right): & -\sigma/4 \le t < \sigma/4 \\ 1: & \sigma/4 \le t < 3\sigma/4 \\ \sin\left(\frac{\pi}{4}\left(5-\frac{4t}{\sigma}\right)\right): & 3\sigma/4 \le t < 5\sigma/4 \\ 0: & \sigma/4 \le t \end{cases}, \quad \omega_k = \frac{\pi}{\sigma}\left(k-\frac{1}{2}\right),$$

for a fixed width σ. The v_k are orthogonal, concentrated on $[0,\sigma]$ (and compactly supported on $[-\sigma/4, 5\sigma/4]$), have a maximum height of $\sqrt{2/\sigma}$, and have at most $3k/2+2$ extrema. These facts combine to give us an upper bound on the total variation:

$$\|v_k(t)\|_{TV} \;\le\; \sqrt{2/\sigma}\,(3k+4).$$

Then Lemma 2 and (19) combine to give us (assuming $K \ge 2$),

$$\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\| \le 2|\theta_1-\theta_2|^{1/2}\left((2/\sigma)\sum_{k=1}^{K}(3k+4)^2\right)^{1/2} \le 12.5\cdot K^{3/2}(|\theta_1-\theta_2|/\sigma)^{1/2},$$

and so we can take $\alpha = 2$ and $N_0 = (12.5)^2K^3/\sigma$, and

$$N(\{\mathcal{S}_\theta\},\epsilon,d) \;\le\; \frac{(12.5)^2K^3T}{\sigma}\epsilon^{-2} \quad\Rightarrow\quad \Delta \le 6\log K + 2\log(T/\sigma) + 16.27.$$

Here we see the dimension K of the subspaces coming out explicitly in the geometric constant.

4 Technical details

4.1 Proof of Theorem 1

Let θ be any member of Θ, and let $\boldsymbol{V}_\theta$ be a $N\times K$ matrix whose columns are an orthobasis for $\mathcal{S}_\theta$. We can write $\boldsymbol{P}_\theta = \boldsymbol{V}_\theta\boldsymbol{V}_\theta^{\mathrm{T}}$ and

$$\|\boldsymbol{P}_\theta - \boldsymbol{P}_\theta\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_\theta\| = \|\boldsymbol{V}_\theta(\mathbf{I}-\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{V}_\theta)\boldsymbol{V}_\theta^{\mathrm{T}}\| = \|\mathbf{I}-\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{V}_\theta\|.$$

By (C1), all the eigenvalues of $\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{V}_\theta$ are between $1\pm\delta_1$, meaning that $(\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{V}_\theta)^{-1}$ exists and is itself close to the identity. We can write

$$\begin{aligned}\tilde{\boldsymbol{P}}_\theta &= \boldsymbol{\Phi}\boldsymbol{V}_\theta(\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{V}_\theta)^{-1}\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}} \\ &= \boldsymbol{\Phi}\boldsymbol{V}_\theta(\mathbf{I}+\boldsymbol{H}_\theta)\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}},\end{aligned}$$

where we can bound the size of $\boldsymbol{H}_\theta$ using a Neumann series:

$$\begin{aligned}\|\boldsymbol{H}_\theta\| &= \left\|(\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{V}_\theta)^{-1}-\mathbf{I}\right\| \\ &= \left\|\sum_{n=1}^{\infty}(\mathbf{I}-\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{V}_\theta)^n\right\| \\ &\le \sum_{n=1}^{\infty}\|\mathbf{I}-\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{V}_\theta\|^n \\ &\le \frac{\delta_1}{1-\delta_1}.\end{aligned}$$

With $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{h}_0$,

$$\begin{aligned}
\left| \|\tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| &= \left| \langle \boldsymbol{y}, \boldsymbol{\Phi}\boldsymbol{V}_\theta(\mathbf{I} + \boldsymbol{H}_\theta)\boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\rangle - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| \\
&\leq \left| \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| + \left| \langle \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}, \boldsymbol{V}_\theta \boldsymbol{H}_\theta \boldsymbol{V}_\theta^{\mathrm{T}}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\rangle \right| \\
&\leq \left| \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| + \|\boldsymbol{V}_\theta \boldsymbol{H}_\theta \boldsymbol{V}_\theta^{\mathrm{T}}\| \, \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\|_2^2 \\
&\leq \left| \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| + \frac{\delta_1}{1-\delta_1} \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\|_2^2. \qquad (23)
\end{aligned}$$

We factor the first term above,

$$\|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 = \left(\|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{h}_0\|_2 + \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2 \right) \left(\|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{h}_0\|_2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2 \right)$$

and then use the fact that

$$\begin{aligned}
\|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{h}_0\|_2 &\leq \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2 + \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_\theta^{\perp} \boldsymbol{h}_0\|_2 \\
&\leq (1+\delta_1)\|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2 + \delta_2,
\end{aligned}$$

to get

$$\left| \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| \leq (2 + \delta_1 + \delta_2)(\delta_1 + \delta_2) =: \delta_3.$$

Returning to (23), we finally have

$$\begin{aligned}
\left| \|\tilde{\boldsymbol{P}}_\theta \boldsymbol{y}\|_2^2 - \|\boldsymbol{P}_\theta \boldsymbol{h}_0\|_2^2 \right| &\leq \delta_3 + \frac{\delta_1(1+\delta_3)}{1-\delta_1} \\
&= \frac{\delta_1 + \delta_3}{1-\delta_1}.
\end{aligned}$$

4.2 Chaining

The following lemma allows us to easily relate a probabilistic bound on the increments of a random process to a probabilistic bound on its supremum. We will use it to establish both Theorems 2 and 3.

Proposition 1 *Let $\{\mathcal{S}_\theta\}$ be a collection of subspaces that is geometrically regular, parameterized by N_0, α as in (12). Let $\boldsymbol{Z}(\theta)$ be a matrix-valued random process indexed by the same parameters $\theta \in \Theta$. If for every $\theta_1, \theta_2 \in \Theta$ we have*

$$\mathrm{P}\left(\|\boldsymbol{Z}(\theta_1) - \boldsymbol{Z}(\theta_2)\| \;\geq\; \gamma(u) \cdot \|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\| \right) \;\leq\; \beta\, e^{-u},$$

where β is a constant and $\gamma(u)$ has the form $\gamma(u) = a\sqrt{u} + bu + cu^2$, then for any fixed $\theta_0 \in \Theta$,

$$\mathrm{P}\left(\sup_{\theta \in \Theta} \|\boldsymbol{Z}(\theta) - \boldsymbol{Z}(\theta_0)\| \geq 3\gamma(u) + 6c \log^2(4^\alpha e) \right) \;\leq\; \beta\, e^{-u+\Delta}, \qquad (24)$$

where $\Delta = \log(8^\alpha N_0^2) + 2$.

If $\boldsymbol{Z}(\theta)$ is vector-valued, then the above holds with the increment size measured in the standard Euclidean norm, $\|\boldsymbol{Z}(\theta_1) - \boldsymbol{Z}(\theta_2)\|_2$.

Proof We use a standard chaining argument (see for example [37, Thm. 1.2.7]).

Let $\mathcal{T}_0 = \{\theta_0\}$, and $\mathcal{T}_j$ be a series of ϵ-approximations of Θ with $\epsilon = 2^{-j}$. We know that we can choose these with $\mathrm{card}(\mathcal{T}_j) \leq N_0 2^{\alpha j}$. For any $\theta \in \Theta$, let $\pi_j(\theta)$ be the closest point in $\mathcal{T}_j$ to θ. Then for every θ, $\|\boldsymbol{Z}(\theta) - \boldsymbol{Z}(\theta_0)\| \leq \sum_{j\geq 0} \|\boldsymbol{Z}(\pi_{j+1}(\theta)) - \boldsymbol{Z}(\pi_j(\theta))\|$. Thus for any given sequence $u_j,\ j = 0, 1, \ldots, \infty$, we have

$$\mathrm{P}\left(\sup_{\theta\in\Theta} \|\boldsymbol{Z}(\theta) - \boldsymbol{Z}(\theta_0)\| \geq \sum_{j\geq 0} \frac{3}{2} 2^{-j}\gamma(u_j)\right) \leq \sum_{j\geq 0} \mathrm{P}\left(\sup_{\theta\in\Theta} \|\boldsymbol{Z}(\pi_{j+1}(\theta)) - \boldsymbol{Z}(\pi_j(\theta))\| \geq \frac{3}{2} 2^{-j}\gamma(u_j)\right).$$

Now since the number of unique pairs of $\pi_{j+1}(\theta), \pi_j(\theta)$ over all θ is at most $\mathrm{card}(\mathcal{T}_j)\,\mathrm{card}(\mathcal{T}_{j+1}) \leq N_0^2 2^{\alpha(2j+1)}$, and $\|\boldsymbol{P}_{\pi_{j+1}(\theta)} - \boldsymbol{P}_{\pi_j(\theta)}\| \leq (3/2)2^{-j}$, we apply the increment tail bound and see that

$$\mathrm{P}\left(\sup_{\theta\in\Theta} \|\boldsymbol{Z}(\theta) - \boldsymbol{Z}(\theta_0)\| \geq \sum_{j\geq 0} \frac{3}{2} 2^{-j}\gamma(u_j)\right) \leq C_a N_0^2 \sum_{j\geq 0} 2^{\alpha(2j+1)} e^{-u_j}.$$

Taking $u_j = u + (2\alpha\log 2 + 1)j + (\alpha \log 2 + 2\log N_0 + 1)$ immediately gives us the righthand side of (24). The lefthand side comes from applying the following simple lemma, which we prove in the appendix.

Lemma 3 *Let $u_j = u + pj + q$, where u, p, and q are positive constants. Let $\gamma(u) = a\sqrt{u} + bu + cu^2$ with $a, b, c \geq 0$. Then*

$$\sum_{j\geq 0} \frac{1}{2} 2^{-j}\gamma(u_j) \leq \gamma(u + p + q) + 2cp^2. \tag{25}$$

■

4.3 Proof of Theorem 2

The proof of Theorem 2 follows from combining Proposition 1, Lemma 4 below, and the fact that for any matrix-valued random process $\boldsymbol{Z}(\theta)$,

$$\sup_{\theta\in\Theta} \|\boldsymbol{Z}(\theta)\| \leq \sup_{\theta\in\Theta} \|\boldsymbol{Z}(\theta) - \boldsymbol{Z}(\theta_0)\| + \|\boldsymbol{Z}(\theta_0)\|, \tag{26}$$

for any fixed $\theta_0 \in \Theta$. For a fixed θ_0, we can bound $\|\boldsymbol{Z}(\theta_0)\|$ using a standard argument; a bound on the first term above is given by Lemma 4 below.

With θ_0 given, the quantity $\|\boldsymbol{P}_{\theta_0} - \boldsymbol{P}_{\theta_0}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_0}\|$ has the exact same distribution as $\|\mathbf{I} - \boldsymbol{A}^{\mathrm{T}}\boldsymbol{A}\|$, where $\boldsymbol{A}$ is a $M \times K$ Gaussian random matrix whose entries have a variance of $1/M$. This random variable describes how far the singular values of an appropriately normalized Gaussian random matrix deviate from 1, and its behavior is well known. In particular, we know (see [42, Lemma 26] which is derived from results in [14]) that

$$\mathrm{P}\left(\|\boldsymbol{P}_{\theta_0} - \boldsymbol{P}_{\theta_0}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_0}\| > 3\gamma(t)\right) \ \leq\ 2e^{-t},$$

for

$$\gamma(t) = \max(\delta(t), \delta^2(t)), \quad \delta(t) = \sqrt{\frac{K}{M}} + \sqrt{\frac{2t}{M}}.$$

We combine this result with the following lemma, which gives us a bound on $\sup_{\theta\in\Theta} \|\boldsymbol{Z}(\theta) - \boldsymbol{Z}(\theta_0)\|$ in (26).

Lemma 4 *Let $\boldsymbol{G}_\theta = \boldsymbol{P}_\theta - \boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{P}_\theta$. There exists a constant C such that for any fixed $\theta_1, \theta_2 \in \Theta$,*

$$\mathrm{P}\left(\|\boldsymbol{G}_{\theta_1} - \boldsymbol{G}_{\theta_2}\| \geq C\,\gamma(t)\,\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|\right) \;\leq\; 5K\,e^{-t}, \tag{27}$$

where

$$\gamma(t) = \sqrt{\frac{Kt}{M}} + \frac{(K + \log(M/K))t}{M} + \frac{t^2}{M}.$$

After establishing the lemma, we can choose θ_0 arbitrarily. Since $K \geq 1$ and the fact that we always have $M \geq \log^2(M)$, we can replace the middle term above with $Kt/M + t/\sqrt{M}$. Applying the chaining proposition from the previous section adds another t^2/M factor (with a constant which depends on α) to $\gamma(t)$ and a factor of e^{Δ} to the failure probability. This establishes the theorem.

Proof of Lemma 4

We will write the increments of $\boldsymbol{G}_\theta$ as a sum of independent random matrices, and then use the matrix Bernstein inequality (Proposition 4) to establish a tail bound. Use $\boldsymbol{\phi}_m \in \mathbb{R}^N$ to denote the rows of $\boldsymbol{\Phi}$, and let $\boldsymbol{S} = \boldsymbol{P}_{\theta_1} + \boldsymbol{P}_{\theta_2}$, and $\boldsymbol{D} = \boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}$. We have

$$\begin{aligned}
\boldsymbol{G}_{\theta_1} - \boldsymbol{G}_{\theta_2} &= \sum_{m=1}^{M} \frac{1}{M}\boldsymbol{P}_{\theta_1} - \frac{1}{M}\boldsymbol{P}_{\theta_2} - \boldsymbol{P}_{\theta_1}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{P}_{\theta_1} + \boldsymbol{P}_{\theta_2}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{P}_{\theta_2} \\
&= \sum_{m=1}^{M} \frac{1}{M}\boldsymbol{D} - \frac{1}{2}\left(\boldsymbol{S}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{S}\right) \\
&=: \sum_{m=1}^{M} \boldsymbol{X}_m,
\end{aligned}$$

Since $\mathrm{E}[\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}] = M^{-1}\mathbf{I}$, the terms in the sum above are zero mean.

Let $\boldsymbol{P}_{\theta_1,\theta_2}$ be the projection operator onto the $2K$ dimensional subspace $\mathcal{S}_{\theta_1} \cup \mathcal{S}_{\theta_2}$, and let

$$R_m = \|\boldsymbol{P}_{\theta_1,\theta_2}\boldsymbol{\phi}_m\|_2, \quad \boldsymbol{\varphi}_m = \boldsymbol{P}_{\theta_1,\theta_2}\boldsymbol{\phi}_m / R_m.$$

We will find it useful to condition on the event that none of the R_m are too large,

$$\mathcal{E}(\xi) = \left\{\sqrt{M} \max_{m=1,\ldots,M} R_m \leq \xi\right\}.$$

Notice that since the $\boldsymbol{\phi}_m$ are Gaussian, the random scalar R_m and random vector $\boldsymbol{\varphi}_m$ are independent; conditioning on $\mathcal{E}(\xi)$ is the same as conditioning the distribution of each $\boldsymbol{X}_m$ on $\mathcal{E}_m(\xi) = \{\sqrt{M}R_m \leq \xi\}$.

We break (27) into two parts:

$$\mathrm{P}\left(\|\boldsymbol{G}_{\theta_1} - \boldsymbol{G}_{\theta_2}\| \geq C\,\gamma(t)\right) \;\leq\; \mathrm{P}\left(\|\boldsymbol{G}_{\theta_1} - \boldsymbol{G}_{\theta_2}\| \geq C\,\gamma(t) \;\mid\; \mathcal{E}_\xi\right) + \mathrm{P}\left(\mathcal{E}(\xi)^c\right). \tag{28}$$

We will bound the first term in (28) using Proposition 4; the second term is subsequently controlled using Proposition 2 (see (30) below).

We can rewrite the $\boldsymbol{X}_m$ as

$$\boldsymbol{X}_m = \frac{1}{M}\boldsymbol{D} - \frac{R_m^2}{2}\left(\boldsymbol{S}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{S}\right).$$

Conditioned on $\mathcal{E}(\xi)$, an upper bound B in Proposition 4 is given by

$$\|\boldsymbol{X}_m\| \leq \frac{\|\boldsymbol{D}\|}{M} + \frac{\xi^2}{M}\|\boldsymbol{S}\boldsymbol{\varphi}_m\|_2\|\boldsymbol{D}\boldsymbol{\varphi}_m\|_2 \leq \frac{2\xi^2+1}{M}\|\boldsymbol{D}\| =: B, \tag{29}$$

since by definition $\|\boldsymbol{\varphi}\|_2 = 1$.

For the conditional variance, since $\mathrm{E}[\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}] = \frac{1}{2K}\boldsymbol{P}_{\theta_1,\theta_2}$,

$$\mathrm{E}[\boldsymbol{S}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{S}] = \frac{1}{K}\boldsymbol{D},$$

and so

$$\mathrm{E}[\boldsymbol{X}_m^2|\mathcal{E}_m(\xi)] = \left(\frac{1}{M^2} - \frac{\mathrm{E}[R_m^2|\mathcal{E}_m(\xi)]}{MK}\right)\boldsymbol{D}^2 + \frac{\mathrm{E}[R_m^4|\mathcal{E}_m(\xi)]}{4}\,\mathrm{E}\left[\left(\boldsymbol{S}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{S}\right)^2\right].$$

We bound this using the fact that $\mathrm{E}[R_m^4|\mathcal{E}_m(\xi)] \leq \mathrm{E}[R_m^4] = 4K(K+1)/M^2$,

$$\begin{aligned}\left\|\mathrm{E}[\boldsymbol{X}_m^2 \mid \mathcal{E}_m(\xi)]\right\| &\leq \frac{\|\boldsymbol{D}\|^2}{M^2} + \frac{\mathrm{E}[R_m^4|\mathcal{E}_m(\xi)]}{4}\left\|\mathrm{E}\left[\left(\boldsymbol{S}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{S}\right)^2\right]\right\| \\ &\leq \frac{\|\boldsymbol{D}\|^2}{M^2} + \frac{K(K+1)}{M^2}\left\|\mathrm{E}\left[\left(\boldsymbol{S}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{S}\right)^2\right]\right\|.\end{aligned}$$

Since R_m and $\boldsymbol{\varphi}_m$ are independent,

$$\mathrm{E}\left[\left(\boldsymbol{S}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\varphi}_m\boldsymbol{\varphi}_m^{\mathrm{T}}\boldsymbol{S}\right)^2\right] = \frac{M^2}{4K(K+1)}\,\mathrm{E}\left[\left(\boldsymbol{S}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{S}\right)^2\right].$$

Expanding the square and making repeated use of Lemma 6 along with the facts,

$$\begin{aligned}\|\boldsymbol{S}\| = \|\boldsymbol{P}_{\theta_1} + \boldsymbol{P}_{\theta_2}\| &\leq 2, \\ \mathrm{trace}(\boldsymbol{D}\boldsymbol{S}) \leq \|\boldsymbol{D}\|_F\|\boldsymbol{S}\|_F &\leq \sqrt{2K}\|\boldsymbol{D}\|\sqrt{2K}\|\boldsymbol{S}\| \leq 4K\|\boldsymbol{D}\|, \\ \mathrm{trace}(\boldsymbol{D}^2) &\leq 2K\|\boldsymbol{D}\|^2, \\ \mathrm{trace}(\boldsymbol{S}^2) &\leq 8K,\end{aligned}$$

results in

$$\left\|\mathrm{E}\left[\left(\boldsymbol{S}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{S}\right)^2\right]\right\| \;\leq\; \frac{32(K+1)}{M^2}\|\boldsymbol{D}\|^2.$$

Thus

$$\left\|\sum_{m=0}^{M}\mathrm{E}[\boldsymbol{X}_m^2 \mid \mathcal{E}_m(\xi)]\right\| \leq \frac{8K+9}{M}\|\boldsymbol{D}\|^2.$$

We now apply Proposition 4 (matrix Bernstein) with the result above along with (29) to get

$$\mathrm{P}\left(\|\boldsymbol{G}_{\theta_1} - \boldsymbol{G}_{\theta_2}\| > C\left(\sqrt{\frac{Kt}{M}} + \frac{\xi^2 t}{M}\right)\|\boldsymbol{D}\| \;\mid\; \mathcal{E}(\xi)\right) \;\leq\; 4Ke^{-t}.$$

Returning to (28), we can combine the bound above along with the following direct consequence of Proposition 2:

$$\mathrm{P}\left(\sqrt{M}\max_{m=1,\ldots,M} R_m > \sqrt{2K} + \tau\right) \;\leq\; Me^{-\tau^2/2}. \tag{30}$$

For and $t, \tau \geq 0$ we have the unconditional bound,

$$\mathrm{P}\left(\|\boldsymbol{G}_{\theta_1} - \boldsymbol{G}_{\theta_2}\| > C\left(\sqrt{\frac{Kt}{M}} + \frac{(\sqrt{2K}+\tau)^2 t}{M}\right)\|\boldsymbol{D}\|\right) \;\leq\; 4Ke^{-t} + Me^{-\tau^2/2}.$$

Taking $\tau = \sqrt{2t + 2\log(M/K)}$ establishes the lemma.

■

4.4 Proof of Theorem 3

We follow the same template as in Section 4.3, but now the process $\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{P}_\theta^{\perp} \boldsymbol{h}_0$ is vector-valued. As before, we have

$$\sup_{\theta\in\Theta} \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{P}_\theta^{\perp} \boldsymbol{h}_0\|_2 \;\leq\; \|\boldsymbol{P}_{\theta_0} \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{P}_{\theta_0}^{\perp} \boldsymbol{h}_0\|_2 \;+\; \sup_{\theta\in\Theta} \|\boldsymbol{P}_\theta \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{P}_\theta^{\perp} \boldsymbol{h}_0 - \boldsymbol{P}_{\theta_0} \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{P}_{\theta_0}^{\perp} \boldsymbol{h}_0\|_2$$

for any fixed θ_0. We can control the first term above using basic properties of Gaussian random vectors; a bound on the second term will follow from Proposition 1 (with $c = 0$) along with Lemma 5 below, which provides a tail bound on the increment of the process.

With θ_0 fixed, the quantity $\boldsymbol{P}_{\theta_0} \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{P}_{\theta_0}^{\perp} \boldsymbol{h}_0$ is a random vector formed by taking a random embedding of $\boldsymbol{h}_0$ followed by an independent random projection. Let $\boldsymbol{V}$ be a $N \times K$ matrix with orthonormal columns such that $\boldsymbol{P}_{\theta_0} = \boldsymbol{V}\boldsymbol{V}^{\mathrm{T}}$, and similarly let $\boldsymbol{U}$ be a $N \times (N-K)$ orthonormal matrix with $\boldsymbol{P}_{\theta_0}^{\perp} = \boldsymbol{U}\boldsymbol{U}^{\mathrm{T}}$. Then

$$\|\boldsymbol{P}_{\theta_0} \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{P}_{\theta_0}^{\perp} \boldsymbol{h}_0\|_2 = \|\boldsymbol{V}^{\mathrm{T}} \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{U} \boldsymbol{g}\|_2, \tag{31}$$

where $\boldsymbol{g} = \boldsymbol{U}^{\mathrm{T}} \boldsymbol{h}_0$; note that $\|\boldsymbol{g}\|_2 \leq \|\boldsymbol{h}_0\|_2 = 1$. Since $\boldsymbol{V}^{\mathrm{T}}\boldsymbol{U} = \boldsymbol{0}$, the $M \times K$ Gaussian random matrix $\boldsymbol{\Phi}\boldsymbol{V}$ and the $M \times (N-K)$ Gaussian random matrix $\boldsymbol{\Phi}\boldsymbol{U}$ are independent of one another. Using a standard concentration bound (Proposition 3 in Section 5), we have

$$\mathrm{P}\left(\|\boldsymbol{V}^{\mathrm{T}} \boldsymbol{\Phi}^{\mathrm{T}} \boldsymbol{\Phi} \boldsymbol{U} \boldsymbol{g}\|_2^2 > \frac{K}{M}\left(1 + \max\left(\frac{7u}{K}, \sqrt{\frac{7u}{K}}\right)\right)\|\boldsymbol{\Phi}\boldsymbol{U}\boldsymbol{g}\|_2\right) \;\leq\; e^{-u},$$

and

$$\mathrm{P}\left(\|\boldsymbol{\Phi}\boldsymbol{U}\boldsymbol{g}\|_2^2 > 1 + \max\left(\frac{7t}{M}, \sqrt{\frac{7t}{M}}\right)\right) \;\leq\; e^{-t}.$$

Taking $u = t$ in (31) yields

$$\mathrm{P}\left(\|\boldsymbol{P}_{\theta_0}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_0}^{\perp}\boldsymbol{h}_0\|_2 > \gamma(t)\right) \;\leq\; 2e^{-t}, \quad \gamma(t) = \sqrt{\frac{K}{M}}\left(1 + \max\left(\frac{7t}{K}, \sqrt{\frac{7t}{K}}\right)\right).$$

For the increment bound, we use the following lemma.

Lemma 5 *Let $\boldsymbol{h}_0 \in \mathbb{R}^N$ be a fixed vector with $\|\boldsymbol{h}_0\|_2 = 1$. There exists a constant C such that for any fixed $\theta_1, \theta_2 \in \Theta$,*

$$\mathrm{P}\left(\|\boldsymbol{P}_{\theta_1}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_1}^{\perp}\boldsymbol{h}_0 - \boldsymbol{P}_{\theta_2}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_2}^{\perp}\boldsymbol{h}_0\|_2 > C\gamma(t)\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|\right) \;\leq\; 4Ke^{-t}, \tag{32}$$

where

$$\gamma(t) = \sqrt{\frac{tK}{M}} + \frac{t\sqrt{K}}{M}.$$

Proof

We can write $\boldsymbol{P}_{\theta_1}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_1}^{\perp}\boldsymbol{h}_0 - \boldsymbol{P}_{\theta_2}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_2}^{\perp}\boldsymbol{h}_0$ as the following sum of independent random vectors,

$$\boldsymbol{P}_{\theta_1}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_1}^{\perp}\boldsymbol{h}_0 - \boldsymbol{P}_{\theta_2}\boldsymbol{\Phi}^{\mathrm{T}}\boldsymbol{\Phi}\boldsymbol{P}_{\theta_2}^{\perp}\boldsymbol{h}_0 \;=\; \sum_{m=1}^{M}(\boldsymbol{P}_{\theta_1}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{P}_{\theta_1}^{\perp} - \boldsymbol{P}_{\theta_2}\boldsymbol{\phi}_m\boldsymbol{\phi}_m^{\mathrm{T}}\boldsymbol{P}_{\theta_2}^{\perp})\boldsymbol{h}_0 \;=:\; \sum_{m=1}^{M}\boldsymbol{x}_m.$$

The $\boldsymbol{x}_m$ above are iid; for convenience, we define

$$\boldsymbol{x} = \frac{1}{2}\left(\boldsymbol{D}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{T}\boldsymbol{h}_0 - \boldsymbol{S}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{h}_0\right), \tag{33}$$

where $\boldsymbol{\phi} \sim \mathrm{Normal}(\mathbf{0}, \mathbf{I})$ and

$$\boldsymbol{D} = \boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2} = \boldsymbol{P}_{\theta_2}^{\perp} - \boldsymbol{P}_{\theta_1}^{\perp}, \quad \boldsymbol{S} = \boldsymbol{P}_{\theta_1} + \boldsymbol{P}_{\theta_2}, \quad \boldsymbol{T} = \boldsymbol{P}_{\theta_1}^{\perp} + \boldsymbol{P}_{\theta_2}^{\perp}.$$

Note that $M\,\mathrm{E}[\|\boldsymbol{x}_m\|_2^2] = \mathrm{E}[\|\boldsymbol{x}\|_2^2]$. We have

$$\|\boldsymbol{x}\|_2^2 = \frac{1}{4}\boldsymbol{h}_0^{\mathrm{T}}\left(\boldsymbol{T}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}^2\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{T} - 2\boldsymbol{T}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{S}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D} + \boldsymbol{D}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{S}^2\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\right)\boldsymbol{h}_0. \tag{34}$$

We will treat each of these three terms separately through applications of Lemma 6. For the first term, we apply

$$\begin{aligned}\mathrm{E}[\boldsymbol{h}_0^{\mathrm{T}}\boldsymbol{T}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}^2\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{T}\boldsymbol{h}_0] &= \boldsymbol{h}_0^{\mathrm{T}}\boldsymbol{T}(2\boldsymbol{D}^2 + \mathrm{trace}(\boldsymbol{D}^2)\mathbf{I})\boldsymbol{T}\boldsymbol{h}_0 \\ &\leq 2\|\boldsymbol{D}\|^2\|\boldsymbol{T}\boldsymbol{h}_0\|_2^2 + \mathrm{trace}(\boldsymbol{D}^2)\|\boldsymbol{T}\boldsymbol{h}_0\|_2^2 \\ &\leq 8(K+1)\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|^2,\end{aligned}$$

since $\|\boldsymbol{T}\boldsymbol{h}_0\|_2^2 \leq 4\|\boldsymbol{h}\|_2^2 = 4$. For the second term in (34),

$$\begin{aligned}\left|\mathrm{E}[\boldsymbol{h}_0^{\mathrm{T}}\boldsymbol{T}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{S}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{h}_0]\right| &= \left|\boldsymbol{h}_0^{\mathrm{T}}\boldsymbol{T}(\boldsymbol{D}\boldsymbol{S} + \boldsymbol{S}\boldsymbol{D} + \mathrm{trace}(\boldsymbol{D}\boldsymbol{S})\mathbf{I})\boldsymbol{D}\boldsymbol{h}_0\right| \\ &\leq (2\|\boldsymbol{S}\|\,\|\boldsymbol{D}\| + \mathrm{trace}(\boldsymbol{D}\boldsymbol{S}))\,\|\boldsymbol{T}\boldsymbol{h}_0\|_2\|\boldsymbol{D}\boldsymbol{h}_0\|_2 \\ &\leq 8(K+1)\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|^2.\end{aligned}$$

Finally for the third term in (34),

$$\begin{aligned}\mathrm{E}[\boldsymbol{h}_0^{\mathrm{T}}\boldsymbol{D}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{S}^2\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{h}_0] &= \boldsymbol{h}_0^{\mathrm{T}}\boldsymbol{D}(2\boldsymbol{S}^2+\mathrm{trace}(\boldsymbol{S}^2)\mathbf{I})\boldsymbol{D}\boldsymbol{h}_0 \\ &\le (2\|\boldsymbol{S}\|^2+8K)\|\boldsymbol{D}\boldsymbol{h}_0\|_2^2 \\ &\le 8(K+1)\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|^2.\end{aligned}$$

Collecting these results, we can take

$$\sigma^2 \;=\; \left\|\sum_{m=1}^{M}\mathrm{E}[\boldsymbol{x}_m^{\mathrm{T}}\boldsymbol{x}_m]\right\| \;=\; M\,\mathrm{E}[\|\boldsymbol{x}_m\|_2^2] \;=\; \frac{1}{M}\,\mathrm{E}[\|\boldsymbol{x}\|_2^2] \;\le\; \frac{8(K+1)}{M}\,\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|^2. \tag{35}$$

Next we need to bound the ψ_1 norm of the $\boldsymbol{x}_m$; we start with

$$\|\boldsymbol{x}_m\|_{\psi_1} = \frac{\|\boldsymbol{x}\|_{\psi_1}}{M}.$$

We will bound each part of $\boldsymbol{x}$ in (33) separately. For the first term, we note that $\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{T}\boldsymbol{h}_0$ is a Gaussian random scalar with variance $\|\boldsymbol{T}\boldsymbol{h}_0\|_2^2$, and so

$$\mathrm{P}\left(|\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{T}\boldsymbol{h}_0|>u\right) \;\le\; \exp\left(-\frac{u^2}{2\|\boldsymbol{T}\boldsymbol{h}_0\|_2^2}\right) \;\le\; \exp\left(-\frac{u^2}{8}\right).$$

Also, $\boldsymbol{D}\boldsymbol{\phi}$ is a Gaussian random vector, and so we have the (standard) bound

$$\mathrm{P}\left(\|\boldsymbol{D}\boldsymbol{\phi}\|_2>u\right) \le\; 2\exp\left(-\frac{u^2}{8\|\boldsymbol{D}\|_F^2}\right) \;\le\; 2\exp\left(-\frac{u^2}{16K\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|^2}\right).$$

Thus for any $t,u\ge 0$,

$$\mathrm{P}\left(\|(\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{T}\boldsymbol{h}_0)\boldsymbol{D}\boldsymbol{\phi}\|_2>t\right) \le \exp\left(-\frac{u^2}{8}\right)+2\exp\left(-\frac{t^2}{u^2 16K\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|^2}\right).$$

Taking $u=t^{1/2}(2K\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|^2)^{-1/4}$ above yields

$$\mathrm{P}\left(\|(\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{T}\boldsymbol{h}_0)\boldsymbol{D}\boldsymbol{\phi}\|_2>t\right) \le 3\exp\left(-\frac{t}{8\sqrt{2K}\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|}\right),$$

and so

$$\|(\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{T}\boldsymbol{h}_0)\boldsymbol{D}\boldsymbol{\phi}\|_{\psi_1} \le 32\sqrt{2K}\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|. \tag{36}$$

We use similar computations for the second part of $\boldsymbol{x}$ in (33):

$$\mathrm{P}\left(|\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{h}_0|>u\right) \le \exp\left(-\frac{u^2}{2\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|^2}\right), \quad \mathrm{P}\left(\|\boldsymbol{S}\boldsymbol{\phi}\|_2>u\right) \;\le\; 2\exp\left(-\frac{u^2}{64K}\right),$$

and so

$$\begin{aligned}\mathrm{P}\left(\|(\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{h}_0)\boldsymbol{S}\boldsymbol{\phi}\|_2>t\right) &\le \mathrm{P}\left(|\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{h}_0|>\frac{\sqrt{\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|t}}{2(2K)^{1/4}}\right)+\mathrm{P}\left(\|\boldsymbol{S}\boldsymbol{\phi}\|_2>\frac{2(2K)^{1/4}\sqrt{t}}{\sqrt{\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|}}\right) \\ &\le 3\exp\left(-\frac{t}{8\sqrt{2K}\|\boldsymbol{P}_{\theta_1}-\boldsymbol{P}_{\theta_2}\|}\right),\end{aligned}$$

and so

$$\|(\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{D}\boldsymbol{h}_0)\boldsymbol{S}\boldsymbol{\phi}\|_{\psi_1} \leq 32\sqrt{2K}\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|. \tag{37}$$

Combining (36) and (37) gives us the upper bound

$$B \;=\; \max_m \|\boldsymbol{x}_m\|_{\psi_1} \;=\; \|\boldsymbol{x}\|_{\psi_1} \;\leq\; \frac{64\sqrt{2K}}{M}\,\|\boldsymbol{P}_{\theta_1} - \boldsymbol{P}_{\theta_2}\|.$$

The lemma follows by applying Proposition 4 with this B and σ^2 as in (35). Since the $\boldsymbol{x}_m$ are all in the union of the column spaces of $\boldsymbol{P}_{\theta_1}$ and $\boldsymbol{P}_{\theta_2}$, we can take $K_1 = 2K$ and $K_2 = 1$ in Proposition 4.

■

5 Probability Tools

The following lemma was useful in establishing the increment bounds in Lemmas 4 and 5. It can be established through a straightforward calculation.

Lemma 6 *Let $\boldsymbol{\phi} \in \mathbb{R}^N$ be a Gaussian random vector, $\boldsymbol{\phi} \sim \mathrm{Normal}(0, \mathbf{I})$, and let $\boldsymbol{A}$ be an arbitrary $N \times N$ matrix. Then*

$$\mathrm{E}\left[\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\boldsymbol{A}\boldsymbol{\phi}\boldsymbol{\phi}^{\mathrm{T}}\right] = \boldsymbol{A} + \boldsymbol{A}^{\mathrm{T}} + \mathrm{trace}(\boldsymbol{A})\,\mathbf{I}.$$

We also make use of the following standard concentration results for the norm of a Gaussian random vector. Nice expositions of these facts can be found in (for example) [30, Prop. 2.18] or [40, A.2.1].

Proposition 2 *Let $\boldsymbol{x} \in \mathbb{R}^K$ be distributed $\boldsymbol{x} \sim \mathrm{Normal}(\mathbf{0}, \mathbf{I})$, and let $\boldsymbol{A}$ be a matrix with N columns. Then*

$$\mathrm{P}\left(\|\boldsymbol{A}\boldsymbol{x}\|_2 \geq \|\boldsymbol{A}\|_F + t\right) \;\leq\; \exp\left(-\frac{t^2}{2\|\boldsymbol{A}\|^2}\right).$$

and

$$\mathrm{P}\left(\|\boldsymbol{A}\boldsymbol{x}\|_2 \geq \lambda\right) \;\leq\; 2\exp\left(-\frac{\lambda^2}{8\|\boldsymbol{A}\|_F^2}\right).$$

Closely related is the following concentration result for the norm of a random matrix applied to a fixed vector.

Proposition 3 *Let $\boldsymbol{x} \in \mathbb{R}^N$ be a fixed vector with $\|\boldsymbol{x}\|_2 = 1$, and let $\boldsymbol{\Phi}$ be an $M \times N$ matrix whose entries are standard independent Gaussian random variables. Then*

$$\mathrm{P}\left(\|\boldsymbol{\Phi}\boldsymbol{x}\|_2^2 > M + \max\left(Ct, \sqrt{CMt}\right)\right) \;\leq\; e^{-t},$$

for a constant $C \leq 2/\log(e/2) \approx 6.52$.

Proof The quantity $\|\boldsymbol{\Phi x}\|_2^2$ is a sum of M independent chi-square random variables. A standard Chernoff bound (see, for example, [41, Chapter 1]) gives us

$$\mathrm{P}\left(\|\boldsymbol{\Phi x}\|_2^2 > M(1+\lambda)\right) \leq ((1+\lambda)e^{-\lambda})^{M/2}.$$

The proposition then follows by applying the inequality

$$(1+\lambda)e^{-\lambda} \;\leq\; e^{-\min(\lambda,\lambda^2)\log(e/2)}, \quad \lambda \geq 0,$$

and taking $t = \min(\lambda, \lambda^2)M/C \;\Rightarrow\; M\lambda = \max(Ct, \sqrt{CMt})$ with $C = 2/\log(e/2)$. ∎

The Matrix Bernstein inequality, stated below, gives us a tight way to estimate the size of a sum of independent random matrices. An early result for this type of bound appeared in [1] and was greatly refined in [39].

Proposition 4 (Matrix Bernstein) *Let $\boldsymbol{X}_1, \ldots, \boldsymbol{X}_M$ be independent self-adjoint random $K \times K$ matrices with $\mathrm{E}[\boldsymbol{X}_m] = \mathbf{0}$ and*

$$\|\boldsymbol{X}_m\| \leq B.$$

Then there exists a universal constant $C_B \leq 4$ such that for all $t \geq 0$,

$$\mathrm{P}\left(\left\|\sum_{m=1}^{M} \boldsymbol{X}_m\right\| \geq C_B\gamma(t)\right) \;\leq\; 2Ke^{-t}, \tag{38}$$

where

$$\gamma(t) = \max\left\{\sigma\sqrt{t},\; 2Bt\right\}, \qquad \sigma^2 = \left\|\sum_{m=1}^{M} \mathrm{E}[\boldsymbol{X}_m^2]\right\|. \tag{39}$$

For non-symmetric $K_1 \times K_2$ matrices $\boldsymbol{X}_m$, we have the similar bound

$$\mathrm{P}\left(\left\|\sum_{m=1}^{M} \boldsymbol{X}_m\right\| \geq C_B\gamma(t)\right) \;\leq\; 2(K_1+K_2)e^{-t},$$

for the same $\gamma(t)$ with

$$\sigma^2 = \max\left\{\left\|\sum_{m=1}^{M} \mathrm{E}[\boldsymbol{X}_m^{\mathrm{T}}\boldsymbol{X}_m]\right\|, \left\|\sum_{m=1}^{M} \mathrm{E}[\boldsymbol{X}_m\boldsymbol{X}_m^{\mathrm{T}}]\right\|\right\}.$$

In our proof of Theorem 3, we find it slightly more convenient to use a version of the matrix Bernstein inequality for matrices that have norms which are subexponential (rather than bounded with probability 1). A convenient way to characterize this property is using the Orlicz norm.

Definition 1 *The Orlicz-ψ_1 norm of a positive random variable X is defined as*

$$\|X\|_{\psi_1} = \inf_{u \geq 0}\left\{u \;:\; \mathrm{E}\left[e^{X/u}\right] \leq 2\right\}.$$

For a random vector $\boldsymbol{x}$, we write $\|\boldsymbol{x}\|_{\psi_1}$ for the Orlicz-ψ_1 norm of $\|\boldsymbol{x}\|_2$; for a random matrix $\boldsymbol{X}$, we write $\|\boldsymbol{X}\|_{\psi_1}$ for the Orlicz-ψ_1 norm of its operator norm $\|\boldsymbol{X}\|$.

If X is subexponential,

$$\mathrm{P}\left(X > u\right) \;\leq\; \alpha\, e^{-\beta u}, \quad \text{for some } \alpha, \beta > 0, \tag{40}$$

then it has finite ψ_1 norm. In fact, a straightforward calculation shows that for X obeying (40),

$$\mathrm{E}\left[e^{X/u}\right] \;\leq\; \left(1 + \frac{1}{\beta u - 1}\right) \alpha^{1/\beta u},$$

and so

$$\|X\|_{\psi_1} \;\leq\; \begin{cases} \frac{\alpha+1}{\beta} & \text{for all } \alpha \geq 1, \\ \frac{2\log\alpha+1}{\beta} & \text{for all } \alpha \geq 3.6, \\ \frac{2\log\alpha}{\beta} & \text{for all } \alpha \geq 18. \end{cases} \tag{41}$$

In Proposition 4, we can replace the condition that the $\|\boldsymbol{X}\| \leq B$ with the condition that the ψ_1-norms can be uniformly bounded,

$$\|\boldsymbol{X}_m\|_{\psi_1} \leq B, \quad m = 1, \ldots, M.$$

Then (38) holds with (see [29, Sec. 2.2])

$$\gamma(t) = \max\left\{\sigma \cdot \sqrt{t},\; 2B\log(2B\sqrt{M}/\sigma) \cdot t\right\}, \tag{42}$$

A Proofs of Lemmas 1 and 2

Assume without loss of generality that $\theta_2 \geq \theta_1$, set $\theta_d = \theta_2 - \theta_1$, and notice that

$$\|v(t-\theta_1) - v(t-\theta_2)\|_2^2 = \int_{-\infty}^{\infty} |v(t) - v(t-\theta_d)|^2 dt.$$

Lemma 1 follows directly from the fact that the Fourier transform preserves energy:

$$\begin{aligned}\int_{-\infty}^{\infty} |v(t) - v(t-\theta_d)|^2 dt &= \frac{1}{2\pi} \int_{-\infty}^{\infty} |\hat{v}(\xi)|^2 \, |1 - e^{-j\xi\theta_d}|^2 \, d\xi \\ &\leq \frac{\theta_d^2}{2\pi} \int_{-\infty}^{\infty} |\hat{v}(\xi)|^2 \, |\xi|^2 \, d\xi.\end{aligned}$$

For Lemma 2, we separate the integral into intervals of length θ_d:

$$\begin{aligned}\int_{-\infty}^{\infty} |v(t) - v(t-\theta_d)|^2 dt &= \int_0^{\theta_d} \sum_{k=-\infty}^{\infty} |v(t+k\theta_d) - v(t+(k-1)\theta_d)|^2 \\ &\leq \int_0^{\theta_d} \left(\sum_{k=-\infty}^{\infty} |v(t+k\theta_d) - v(t+(k-1)\theta_d)| \right)^2 \\ &\leq \theta_d \, \|v\|_{TV}^2.\end{aligned}$$

B Proof of Lemma 3

We have

$$\sum_{j\geq 0} \frac{1}{2} 2^{-j} g(u_j) = \sum_{j\geq 0} \frac{1}{2} 2^{-j} \left(a\sqrt{u_j} + b u_j + c u_j^2 \right).$$

Since $a\sqrt{u_j} + b u_j$ is a concave function of u_j and since $\sum_{j\geq 0} 2^{-j-1} = 1$ and also $\sum_{j\geq 0} j 2^{-j-1} = 1$, by the Jensen inequality

$$\begin{aligned}\sum_{j\geq 0} \frac{1}{2} 2^{-j} \left(a\sqrt{u_j} + b u_j \right) &\leq a \sqrt{\sum_{j\geq 0} \frac{1}{2} 2^{-j} u_j} + b \sum_{j\geq 0} \frac{1}{2} 2^{-j} u_j \\ &= a\sqrt{u+p+q} + b(u+p+q). \qquad (43)\end{aligned}$$

Now using the fact that $\sum_{j\geq 0} j^2 2^{-j-1} = 3$,

$$\begin{aligned}\sum_{j\geq 0} \frac{1}{2} 2^{-j} c u_j^2 &= c \sum_{j\geq 0} \frac{1}{2} 2^{-j} (u^2 + p^2 j^2 + q^2 + 2upj + 2uq + 2pqj) \\ &= c(u^2 + 3p^2 + q^2 + 2up + 2uq + 2pq) \\ &= c(u+p+q)^2 + 2cp^2. \qquad (44)\end{aligned}$$

The lemma is proved by combining (43) and (44).

References

[1] R. Ahlswede and A. Winter. Strong converse for identification via quantum channels. *IEEE Trans. Inform. Theory*, 48(3):569–579, 2002.

[2] R. G. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.

[3] R. G. Baraniuk and M. B. Wakin. Random projections of smooth manifolds. *Found. of Comput. Math.*, 9(1):51–77, 2009.

[4] T. Blumensath. Sampling and reconstructing signals from a union of subspaces. *IEEE Trans. Inform. Theory*, 57(7):4660–4671, 2011.

[5] T. Blumensath and M. E. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE Trans. Inform. Theory*, 55(4):1872–1882, April 2009.

[6] P. Boufounos, G. Kutyniok, and H. Rauhut. Sparse recovery from combined fusion frame measurements. *IEEE Trans. Inform. Theory*, 57(6):3864–3876, June 2011.

[7] J. Bourgain and J. Nelson. Toward a unified theory of sparse dimensionality reduction in Euclidean space. arxiv:1311.2542, November 2013.

[8] E. Candès. The restricted isometry property and its implications for compressed sensing. *C. R. Acad. Sci. Paris, Ser. I*, 346:589–592, 2008.

[9] E. Candès and Y. Plan. Tight oracle bounds for low-rank matrix recovery from a minimal number of random measurements. *IEEE Trans. Inform. Theory*, 57(4):2342—2359, 2011.

[10] E. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. on Pure and Applied Math.*, 59(8):1207–1223, 2006.

[11] E. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Trans. Inform. Theory*, 52(12):5406–5245, December 2006.

[12] K. L. Clarkson. Tighter bounds for random projections of manifolds. In *Proc. 24th Ann. Symp. Comput. Geom.*, pages 39–48. ACM, 2008.

[13] S. Dasgupta and A. Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Struct. and Alg.*, 22(1):60–65, 2003.

[14] K. R. Davdison and S. J. Szarek. Local operator theory, random matrices and Banach spaces. In W. B. Johnson and J. Lindenstrauss, editors, *Handbook of the Geometry of Banach Spaces*, volume 1, chapter 8, pages 317–366. Elsevier, 2001.

[15] M. A. Davenport, P. T. Boufounos, M. B. Wakin, and R. G. Baraniuk. Signal processing with compressive measurements. *IEEE J. Selected Topics in Sig. Proc.*, 4(2):445–460, 2010.

[16] M. A. Davenport, M. F. Duarte, M. B. Wakin, J. N. Laska, D. Takhar, K. F. Kelly, and R. G. Baraniuk. The smashed filter for compressive classification and target recognition. In *Proc. SPIE Conference on Computational Imaging V*, San Jose, CA, January 2007.

[17] S. Dirksen. Tail bounds via generic chaining. arxiv:1309.3522, September 2013.

[18] S. Dirksen. Dimensionality reduction with subgaussian matrices: a unified theory. arxiv:1402.3973, February 2014.

[19] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, April 2006.

[20] A. Eftekhari, J. Romberg, and M. B. Wakin. Matched filtering from limited frequency samples. *IEEE Trans. Inform. Theory*, 59(6):3475–3496, 2013.

[21] C. Ekanadham, D. Tranchina, and E. P. Simoncelli. Recovery of sparse translation-invariant signals with continuous basis pursuit. *IEEE Trans. Sig. Proc.*, 59(10):4735–4744, October 2011.

[22] Y. C. Eldar, P. Kuppinger, and H. Bolcskei. Block-sparse signals: Uncertainty relations and efficient recovery. *IEEE Trans. Sig. Proc.*, 58(6):3042–3054, June 2010.

[23] Y. C. Eldar and M. Mishali. Robust recovery of signals from a structured union of subspaces. *IEEE Trans. Inform. Theory*, 55(11):5302–5316, November 2009.

[24] K. Fhyn, M. F. Duarte, and S. H. Jensen. Compressive parameter estimation for sparse translation-invariant signals using polar interpolation. arxiv:1305.3483, September 2013.

[25] K. Gedalyahu and Y. C. Eldar. Time-delay estimation from low-rate samples: a union of subspaces approach. *IEEE Trans. Sig. Proc.*, 58(6):3017–3031, June 2010.

[26] Finn B. Jensen, William A. Kuperman, M. B. Porter, and Henrik Schmidt. *Computational Ocean Acoustics*. AIP Press, 2000.

[27] W. B. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Conf. in Modern Analysis and Probability*, pages 189–206, 1984.

[28] B. Klartag and S. Mendelson. Empirical processes and random projections. *J. Funct. Anal.*, 225(1):229–245, 2005.

[29] V. Koltchinkskii. Von Neumann entropy penalization and low-rank matrix completion. *Ann. Statist.*, 39(6):2936–2973, 2012.

[30] M. Ledoux. *The concentration of measure phenomenon*, volume 89. Amer Mathematical Society, 2001.

[31] H. S. Malvar and D. H. Staelin. The LOT: Transform coding without blocking effects. *IEEE Trans. Acoustics, Speech, and Signal Proc.*, 37:553–559, April 1989.

[32] W. Mantzel, J. Romberg, and K. Sabra. Compressive Matched-Field Processing. *J. Acoust. Soc. of Amer.*, 132:90–102, 2012.

[33] V. A. Marchenko and L. A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Mat. Sb. (N.S.)*, 72(114)(4):507–526, 1967.

[34] S. Mendelson. Empirical processes with a bounded ψ_1 diameter. *Geom. Funct. Anal.*, 20:988–1027, 2010.

[35] S. Mendelson, A. Pajor, and N. Tomczak-Jaegermann. Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geometric and Functional Analysis*, 17(4):1248–1282, 2007.

[36] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Review*, 52(3):471–501, 2010.

[37] M. Talagrand. *The Generic Chaining: Upper and Lower Bounds on Stochastic Processes*. Springer, 2005.

[38] G. Tang, B. N. Bhaskar, and B. Recht. Near minimax line spectral estimation. arxiv:1303.4348, March 2013.

[39] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Found. of Comput. Math.*, 12(4):389–434, 2012.

[40] A. W. van der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes*. Springer, 1996.

[41] S. Vempala. *The Random Projection Method*, volume 65 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. AMS, 2004.

[42] R. Vershynin. Introduction to the non-asymptotic theory of random matrices. In Y. Eldar and G. Kutyniok, editors, *Compressed Sensing, Theory and Applications*, pages 210–268. Cambridge University Press, 2012.

[43] H. L. Yap, M. B. Wakin, and C. J. Rozell. Stable manifold embedding with structured random matrices. *J. Sel. Topics Sig. Proc.*, 7(4):720–730, 2013.

[44] J. Yoo, C. Turnes, E. Nakamura, C. Le, S. Becker, E. Sovero, M. Wakin, M. Grant, J. Romberg, A. Emami-Neyestanak, and E. Candès. A compressed sensing parameter extraction platform for radar pulse signal acquisition. *IEEE J. Emerging Topics Cir. and Sys.*, 2(3):626–638, 2012.